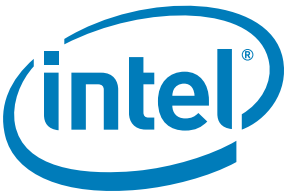

Intel® 631xESB/632xESB I/O Controller Hub

Specification Update

January 2008

Document number: 313075-006

Contents

§

Revision History

Version	Description	Date
001	• Initial Release	May 2006
002	• Updated Component Marketing Information Table • Change code name ESB2 to product name Intel® 631xESB/ 632xESB I/O Controller Hub • Added Errata 96 • Added TPM support description • Other Documentation Changes	August 2006
003	• Added Errata 97-98 • Added Documentation Changes	February 2007
004	• Added S-Spec information • Added Erratum 99 • Added Documentation Changes item 12-14	April 2007
005	• Added Errata 100 • Added Documentation Changes 10-13 and renumbered ordering	July 2007
006	• Added Erratum 101 • Added Documentation Change 19	January 2008

1 Preface

This is an update to the specifications in the documents listed in the "Affected Documents" and "Related Documents" tables. It is a compilation of device and document errata and specification clarifications/changes, and is intended for hardware system manufacturers and software developers.

Information types defined in the Nomenclature section of this document are consolidated into this document and are no longer published in other documents. This document may also contain previously unpublished information.

1.1 Affected Documents

Document Title	Document Number
Intel® 631xESB/632xESB I/O Controller Hub EDS vols 1-3 rev 2.0	20850, 20852, 20853

1.2 Related Documents

Document Title	Document Number
IS-Bensley/Bensley-VS Platform Design Guide (PDG) rev 1.5	20757
RS-Intel® Enterprise South Bridge 2 (ESB2) BIOS Specification rev 0.7	19052
Intel® Enterprise South Bridge 2 (ESB2) BIOS Specification Update rev 0.77	iBL

1.3 Nomenclature

S-Spec Number is a five-digit code used to identify products. Products are differentiated by their unique characteristics,e.g., core speed, L3 cache size, package type, etc. as described in the processor identification information table. Read all notes associated with each S-Spec number.

Errata are design defects or errors. These may cause the ESB2 behavior to deviate from published specifications. Hardware and software designed to be used with any given stepping must assume that all errata documented for that stepping are present on all devices.

Specification Changes/Clarifications are modifications to the current published specifications. These changes will be incorporated in the next release of the specification.

Documentation Changes include typos, errors, or omissions from the current published specifications. These will be incorporated in the next release of the specification.

Note: Errata remain in the specification update throughout the product's life cycle, or until a particular stepping is no longer commercially available. Under these circumstances, errata removed from the specification update are archived and available upon request. Specification changes, specification clarifications, and documentation changes are

removed from the specification update when the appropriate changes are made to the appropriate product specification or user documentation.

2 Summary Table of Changes

The tables included in this section indicate the errata, specification changes/ clarifications, or documentation changes that apply to the Intel® 631xESB/632xESB I/O Controller Hub. Intel may fix some of the errata in a future stepping of the component, and account for the other outstanding issues through documentation or specification changes as noted.

2.1 Codes Used in Summary Table

Stepping / Version

X: Applies to this stepping.

Blank: Fixed in listed stepping or does not exist in listed stepping.

Status

No Fix - Root caused to a silicon issue that will not be fixed.

Plan Fix - Root cased to a silicon issue and will be fixed in a future stepping.

Fixed - Root cased to a silicon issue and has been fixed in a subsequent stepping.

Spec Change - Root caused to a specification error that will be updated.

Spec Clarification - Root caused to a specification requiring clarification

Document Update - Root caused to a documentation error that will be updated

Row

Change bar to left of table row indicates that this item is either new or modified from the previous version of this document.

2.2 Errata

Number	Intel® 631xESB/ 632xESB I/ O Controller Hub Stepping		Status	SKUs (optional)	ERRATA
	A-0	A-1			
1	x		Fixed	All	Unable to write to byte 2 of PCIXPM_PMCSR register when Byte Enables (BE) not 0xF
2	x	x	No Fix	All	Poisoned PCI Express* transaction (TLP) causes error to be logged in PCIX Bridge unit
3	x	x	No Fix	All	Intel® 631xESB/632xESB I/O Controller Hub not logging header for PCI Express* unsupported requests when detected at receiver
4	x	x	No Fix	All	PCI Express* Replay timer is not frozen in Recover State
5	x	x	No Fix	All	PCI Express* 2 ms “fail-safe” timer to exit from hot-reset doesn’t work
6	x	x	No Fix	All	Some inbound PCI Express* transactions that require master abort are logged incorrectly
7	x	x	No Fix	All	Default value used to calculate L0s exit latency is incorrect.
8	x		Fixed	All	Wrong and illegal negotiated width value can be reported for downstream PCI Express* ports
9	x	x	No Fix	All	HPCCTL register incorrectly located in configuration space
10	x		Fixed	All	Max Payload Size register field has wrong default
11	x	x	No Fix	All	Cache Line Size register should be type RW
12	x		Fixed	All	SATA AHCI version register value incorrect
13	x	x	No Fix	All	SATA EB buffer overflow should set ERR.E bit instead of ERR.M bit.
14	x		Fixed	All	Output of the JTAG IDCODE register is a constant ‘1’
15	x		Fixed	All	Upstream Port captures incorrect value in Device Capabilities register when receives Set_slot_pwr Message
16	x		Fixed	All	Incorrect SET_SLOT_PWR Message data payload on Intel® 631xESB/ 632xESB I/O Controller Hub external downstream ports
17	x		Fixed	All	Reading LAN PCI function BARS can return incorrect values
18	x		Fixed	All	BMC acting as Fast Management Link Master, asserting the FLBINTEX signal can cause loss of synchronization between master and slave
19	x		Fixed	All	Broadcast Packets Transmitted Counter may wrongly count flow control packets
20	x	x	No Fix	All	Good Packets Received Counter may be incorrect due to Link disconnect
21	x	x	No Fix	All	Intel® 631xESB/632xESB I/O Controller Hub LAN: FML hold time violation per specification
22	x	x	No Fix	All	Packets with Symbol error may be counted as good packets
23	x		Fixed	All	RNBC counter is not correct
24	x		Fixed	All	gtx_clk is driven from the crystal in MII 100/10 mode is not synchronized to 125 MHz internal clock.
25	x		Fixed	All	ESI link has no de-emphasis
26	x	x	No Fix	All	SATA Default Tx drive strength set too high at Gen1/Gen2.
27	x		Fixed	All	Test port pins that are used for straps are not sampled correctly at PE RESET
28	x	x	No Fix	All	PATA config space disappears when secondary IO decode is turned on in native mode.
29	x		Fixed	All	PCI Express* lanes not driving to correct level.

2.2 Errata (Continued)

Number	Intel® 631xESB/632xESB I/O Controller Hub Stepping		Status	SKUs (optional)	ERRATA
	A-0	A-1			
30	x		Fixed	All	IDE setting SRST (Software Reset) doesn't set the BSY bit
31	x		Fixed	All	MAC Transmission of Preamble violates 802.3 specification
32	x	x	No Fix	All	SATA Gen2 Driver current lower than expected
33	x	x	No Fix	All	PCI Express* Link Control Register requires word write to update.
34	x		Fixed	All	PCI Express* downstream ports flag inbound 4K memory read as malformed Transaction Layer Packet (TLP)
35	x		Fixed	All	PCI Express* Hot-Plug: HPCCTL register not sticky
36	x		Fixed	All	PCI Express* Hot-Plug: Start-of-day Power Fault not latched for port B in two single-byte mode.
37	x		Fixed	All	Intel® 631xESB/632xESB I/O Controller Hub unable to enter the PCI Express* Power Management Link States L0s, L1 and L2 on ports PE1, PE2, and Intel® 631xESB/632xESB I/O Controller Hub LAN.
38	x		Fixed	All	PCI Express* Hot-Plug: Hot swap to lower link width fails.
39	x	x	No Fix	All	PCI Express* port PE0 Link layer should drop Data Link Layer Packets (DLLPs) with unknown encoding type.
40	x	x	No Fix	All	Intel® 631xESB/632xESB I/O Controller Hub port PE0 does not ignore a PCI Express* Null Packet.
41	x	x	No Fix	All	Intel® 631xESB/632xESB I/O Controller Hub sending less than the minimum number of Power Management Acknowledgments (PMAKs) to SATA.
42	x	x	No Fix	All	Advanced Host Controller Interface (AHCI): Improper length register device-to-host FIS.
43	x		Fixed	All	Standard Hot-Plug Controller (SHPC) Message Signaled Interrupt (MSI) is lost/corrupted after PCI config access to 0x78 or 0x7C.
44	x	x	No Fix	All	PCI Express* Extended Tag Capability Bit
45	x	x	No Fix	All	Command register Bus Number field attribute not PCI compliant for the PCI-X bridge
46	x		Fixed	All	PxSERR errors are logged during BIOS initialization when running at 3Gbps on a small percentage of parts
47	x		Fixed	All	Intel® 631xESB/632xESB I/O Controller Hub LAN:When the IDE disabled by FW all the BAR's return 7F
48	x	x	No Fix	All	PCI Express* Link Training and Status State Machine (LTSSM) on ports PE1and PE2 incorrectly sends TS1-PAD-PAD in Configuration.Linkwidth.Start
49	x		Fixed	All	Non-Hot Plug Presence Detect bit not set on PCI Express* ports PE1 and PE2.
50	x	x	No Fix	All	Intel® 631xESB/632xESB I/O Controller Hub LAN not handling Lock Transactions as Unsupported Requests per PCI Express* spec.
51	x	x	No Fix	All	Downstream PCI Express* bridges are inaccessible if secondary bus number misconfigured
52	x	x	No Fix	All	Advanced Host Controller Interface (AHCI) Host Bus Adapter (HBA) BSY bit set when recovering from fatal error
53	x	x	No Fix	All	Noise on PCI Express* TX coming out of Electrical Idle
54	x	x	No Fix	All	IDE VIL not meeting spec
55	x	x	No Fix	All	JTAG read data corrupted when concurrent config access
56	x		Fixed	All	Standard Hot-Plug Controller (SHPC) Hot-Plug PM_PME Requestor ID not correct.

2.2 Errata (Continued)

Number	Intel® 631xESB/ 632xESB I/O Controller Hub Stepping		Status	SKUs (optional)	ERRATA
	A-0	A-1			
57	x	x	No Fix	All	PCI Express* Completion timer not halting in L1
58	x	x	No Fix	All	Split lock cycle to LPC space result in FSB timeout and IERR
59	x	x	No Fix	All	PCI Express* SKP/InitFCx Contention
60	x	x	No Fix	All	Upstream PCI Express* Ports (PE3 & PE4) enter the L1 link state when all downstream devices are NOT in the d3hot power management state
61	x	x	Closed	All	Interrupt Disable attribute does not correspond to implementation for upstream PCI-E bridge
62	x	x	Closed	All	Max Read Request Size config default not PCI compliant for all PCI-E ports
63	x	x	No Fix	All	Under certain conditions, inbound prefetched PCI read requests may return wrong data to the requestor
64	x	x	No Fix	All	USB 2.0 Transmit Ring-Back Failure.
65	x	x	No Fix	All	Intel® 631xESB/632xESB I/O Controller Hub LAN: Cannot read from EEPROM by parallel access on some EEPROMS
66	x	x	No Fix	All	PCI-X clock falling edge slew rate fails spec into test load.
67	x	x	No Fix	All	Flow Control Completion Data credits may incorrectly update based on length field
68	x	x	No Fix	All	Intel® 631xESB/632xESB I/O Controller Hub LAN: Missed packets in 10/ 100Mb HD
69	x	x	No Fix	All	Poisoned TLP generated on PCI Express* MSI hot plug interrupt
70	x	x	No Fix	All	Intel® 631xESB/632xESB I/O Controller Hub LAN: Certain values of last ip_option may corrupt RSS decision
71		x	No Fix	All	PCI EXpress* Hot-Plug interrupt does not send simultaneous with PM_PME
72	x	x	No Fix	All	SATA Protocol errors detected in PxSERR register during boot
73	x	x	No Fix	All	Intel® 631xESB/632xESB I/O Controller Hub LAN: Invalid commas on the Serdes interface corrupts the data integrity of the preceding packet
74	x	x	No Fix	All	SATA BSY timeout after Standby Immediate and SRST (after 1ms)
75	x	x	No Fix	All	Unsolicited COMINIT while FIS posting pending will corrupt the FIS posting cycle
76	x	x	No Fix	All	Link Down at Upstream PCI Express* is resulting in failure of training on the Downstream PCI Express*
77	x	x	No Fix	All	Link Control Retrain Link bit not reading 0 when slot is empty
78	x	x	No Fix	All	Intel® 631xESB/632xESB I/O Controller Hub LAN: Return Loss in SERDES
79	x	x	No Fix	All	A PCI Express Downstream ports detect correctable errors while running L0s
80	x	x	No Fix	All	PCI Express Hot-Plug disable when downstream link is in L0s/L1 can hang the system
81	x	x	No Fix	All	Intel® 631xESB/632xESB I/O Controller Hub LAN: Link fail in middle of Rx packet may corrupt data
82	x	x	No Fix	All	Intel® 631xESB/632xESB I/O Controller Hub LAN: Does not resent pause packet
83	x	x	No Fix	All	Intel® 631xESB/632xESB I/O Controller Hub LAN: Ethernet frame length issues
84	x	x	No Fix	All	Intel® 631xESB/632xESB I/O Controller Hub LAN: IDE BAR4 offset2
85	x	x	No Fix	All	Intel® 631xESB/632xESB I/O Controller Hub LAN: FML - Clock extending on the stop cycle

2.2 Errata (Continued)

Number	Intel® 631xESB/ 632xESB I/O Controller Hub Stepping		Status	SKUs (optional)	ERRATA
	A-0	A-1			
86	x	x	No Fix	All	Intel® 631xESB/632xESB I/O Controller Hub LAN: Serdes is unable to acquire synchronization from ordered sets beginning with K28.1 and K28.7
87	x	x	No Fix	All	Intel® 631xESB/632xESB I/O Controller Hub LAN: Formed and invalid /C/ code handling
88	x	x	No Fix	All	Intel® 631xESB/632xESB I/O Controller Hub LAN: False detection on an idle_match condition
89	x	x	No Fix	All	Intel® 631xESB/632xESB I/O Controller Hub LAN: Ability Match and Acknowledge Match
90	x	x	No Fix	All	Intel® 631xESB/632xESB I/O Controller Hub LAN: ROM FW - Bus may hang when Master reads more bytes than Slave reported
91	x	x	No Fix	All	Intel® 631xESB/632xESB I/O Controller Hub LAN: ROM FW - SOL Timeout character control byte in EEPROM image
92	x	x	No Fix	All	Intel® 631xESB/632xESB I/O Controller Hub LAN: ROM FW - BMC fragments that are sent through 2 different SMBus ports are sent over LAN as a single packet
93	x	x	No Fix	All	Intel® 631xESB/632xESB I/O Controller Hub LAN - Time until Jam is longer than it should be
94	x	x	No Fix	All	Configuration cycle failure when targeting PCI Express* PE1 and PE2 while PM L1 entry sequence is in process
95	x	x	No Fix	All	PCI Express* ports may gate PM L1 entry based on lack of Flow Control credit availability
96	x	x	No Fix	All	ESB2 SATA Signal Voltage Level
97	x	x	No Fix	All	Link Control register bit 3 in PCI Express Capability table should be read-only and always return 0 for switch ports
98	x	x	No Fix	All	Device Control register bit 8 should be read-writable if Extended Tag Field Support is indicated
99	x	x	No Fix	All	CPU IERR during POST and CPU IERR with hang at OS prompt
100	x	x	No Fix	All	"ESB2 PCIe switch mis-formats TLPs with ECRC destined to the MCH"
101	x	x	No Fix	All	Enabling RSS in the Middle of Received Packets May Stop Receive Flow

2.3 Specification Changes

Number	SPECIFICATION CHANGES
1	N/A

2.4 Specification Clarifications

Number	SPECIFICATION CHANGES
1	PCI Downstream Device Disable Clarification

2.5 Documentation Changes

Number	DOCUMENTATION CHANGES
1	Content updated to Section 25.1.34
2	Content updated to Section 25.2.34
3	Content updated to Section 25.3.32
4	Content updated to Section 25.4.33
5	Content updated to Section 25.5.35
6	Content updated to Section 25.6.33
7	Content updated to Table4-1
8	Content added to Section 5.25
9	Content modified to section 23.1.34
10	Content updated to Table6-1 and Table6-2
11	“Content updated to Table6-8”
12	Content updated to Table2-32
13	“Content updated to Section13.6.1.64”
14	“Content updated to Section 25.1.34”
15	Added Table6-20 Intel® 631xESB/632xESB I/O Controller Hub Absolute Maximum Ratings
16	Section 8.1, changed pin AA36 name from “Reserved” to “PHY_POWER_DOWN”, changed pin AB36 name from “Reserved” to “PHYRST#”
17	“Section 8.2, changed pin AA36 name from “Reserved” to “PHY_POWER_DOWN”, changed pin AB36 name from “Reserved” to “PHYRST#””
18	“Input Signal Behavior Clarifications”
19	APM_CNT and APM_STS Registers

3 Identification Information

3.1 Component Identification via Programming Interface

The Intel® 631xESB/632xESB I/O Controller Hub can be identified by the following register contents:

Stepping	Vendor ID[a]	Device ID[b]	Revision Number[c]
A-0	8086h	244Eh	D8h
A-1	8086h	244Eh	D9h

Notes:

a. The Vendor ID corresponds to bits 15:0 of the Vendor ID Register located at offset 00 - 01h in the PCI function 0 configuration space.

b. The Device ID corresponds to bits 15:0 of the Device ID Register located at offset 02 - 03h in the PCI function 0 configuration space.

c. The Revision Number corresponds to bits 7:0 of the Revision ID Register located at offset 08h in the PCI function 0 configuration space.

3.2 Component Marking Information

The Intel® 631xESB/632xESB I/O Controller Hub stepping can be identified by the following component markings:

Stepping	QDF-Spec	SKU	Top Marking	Notes
A-1, leaded	SL97P	Intel® 6321ESB I/O Controller Hub (ESB2-E)	NQ80003ES2	Production
A-1, Pb free	SL97Q	Intel® 6321ESB I/O Controller Hub (ESB2-E)	QG80003ES2	Production
A-1, leaded	SL97M	Intel® 6311ESB I/O Controller Hub (ESB2-T)	NQ80003ES2	Production
A-1, Pb free	SL97N	Intel® 6311ESB I/O Controller Hub (ESB2-T)	QG80003ES2	Production

Figure 3-1. Top-Side Marking Example

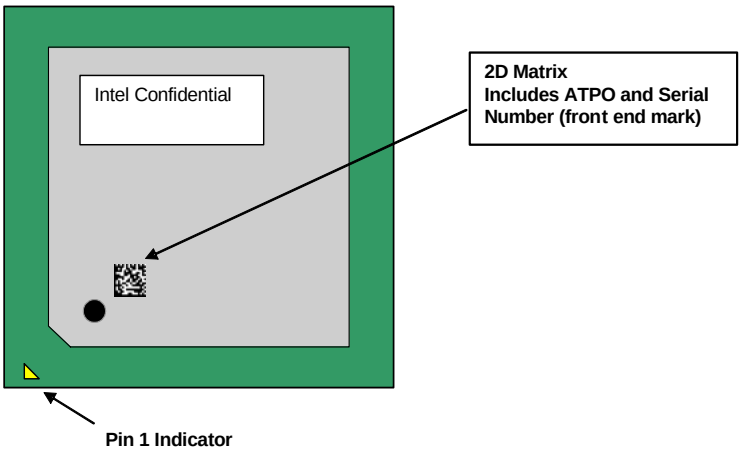

Figure 3-2. Top-Side Photo

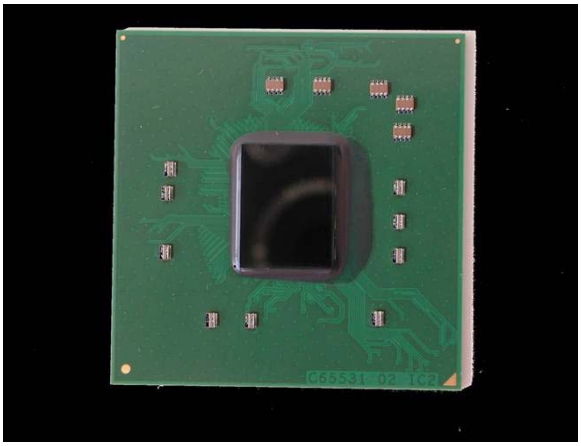

Figure 3-3. Bottom-Side Photo

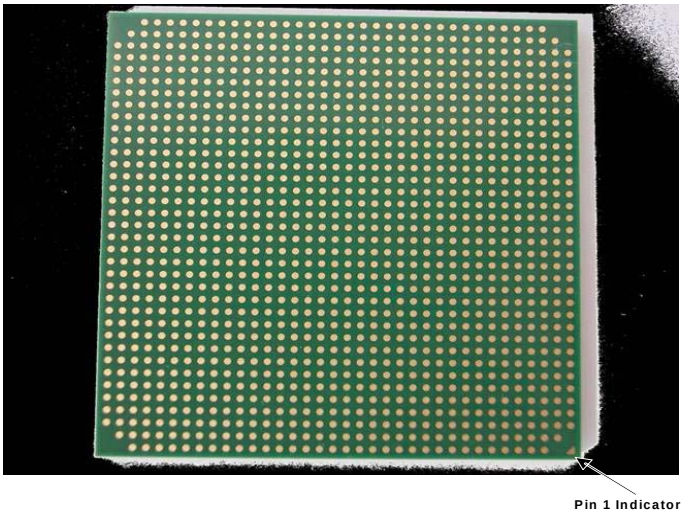

4 Errata

1. Unable to write to byte 2 of PCIXPM_PMCSR register when Byte Enables (BE) not 0xF

Problem: Register PCIXPM_PMCSR (Bm:D0:F3 offset 0x70) in the upstream function control registers (PE3 and PE4) requires full word reads or writes. Partial word reads or writes return 0s.

Implication: Failure was encountered while running the power management capability test provided in PCI-SIG test suite.

Workaround: Software may try to write to clear the status bit without affecting D state. To do this (given this bug), a read of register would need to happen first before the write to ensure nothing affects the Dstate.

Status: Fixed

2. Poisoned PCI Express* Transaction Layer Packet (TLP) causes error to be logged in PCIX Bridge

Problem: Configuration transactions (both type 0 and type 1) with poisoned TLP targeting either downstream PCI Express* port or upstream PCI Express* port (ports PE1 through PE4) cause error to be logged incorrectly. The error is logged in the PCIX Bridge interface (function 3); not PCI Express* interfaces as expected. Since the cycle is poisoned (EP bit set), the cycle will already be logged at the source of the transaction.

Implication: Some error handling routines might get confused.

Workaround: Error handling software must comprehend this issue.

Status: No Fix

3. Downstream PCI Express* interface not logging header for PCI Express* unsupported requests when detected at receiver

Problem: Downstream PCI Express* interface (PE1 and PE2) is not logging the header for unsupported request (URREQ) transactions detected at receiver side. Certain illegal inbound transactions are not fully logged as “unsupported requests.” While the status bits are set correctly and error messages are generated correctly, the header log is not set. The following illegal transactions are affected: Vendor Defined Type 0 Message routed by ID targeting Intel® 631xESB/632xESB I/O Controller Hub internal bus. PM active state NAK message with TC = 000. Vendor Defined Messages with routing set to “Terminate at receiver. Messages with routing “terminate at receiver” – with encodings not defined in Spec. Receive Set_Slot_Power_Limit message inbound with tc = 000. Messages except PME_TO_ACK with routing as “Gather and route to RC. Messages with reserved routing fields

Implication: This bug violates the PCI Express* Specification. It is also in the compliance checklist: SYS.2.7#20 - If an unsupported request is received the receiver must log the header of the TLP that caused the error in the header log register if the register is currently invalid. This bug violates the PCI Express* Specification but is not expected to cause any functional failures.

Workaround: NA

Status: No Fix

4. PCI Express* Replay timer is not frozen in Recover State

Problem: The PCI Express* Replay Timer is a fail-safe timer defined in the PCI Express* Specification such that if a link does not receive an Ack or Nack within a predefined

time, it assumes it received a Nack and resends the packet(s). According to the PCI Express* spec (section 3.5.2.1) the Replay Timer should be frozen during link retraining. The bug is that in Intel® 631xESB/632xESB I/O Controller Hub, the Replay Timer continues counting during link retraining. This means Intel® 631xESB/632xESB I/O Controller Hub can retry a transaction after retraining sooner than it should (bug exists in PE1 though PE4 PCI Express* ports).

Implication: This bug violates the PCI Express* Specification but is not expected to cause any functional failures.

Workaround: NA

Status: No Fix

5. PCI Express* 2 ms “fail-safe” timer to exit from hot-reset doesn’t work

Problem: The 2 ms timeout signal transition from Hot-reset to Detect doesn’t occur when LTSSM is in hot reset. When a PCI Express* link is put in hot reset, the specification defines a 2 ms timeout timer as a “fail-safe” exit mode from hot reset. See PCI Express* 1.0a specification section 4.2.6.10, page 199: “Otherwise, after a 2 ms timeout next state is Detect”. In Intel® 631xESB/632xESB I/O Controller Hub, the timer doesn’t work. The upstream PCI Express* port (PE4) correctly exits hot reset when the hot reset condition is removed and the bus goes into electrical idle. The downstream PCI Express* ports (PE1 and PE2) correctly exit hot reset when warm reset goes inactive.

Implication: This bug violates the PCI Express* Specification but is not expected to cause any functional failures. The 2 ms is a fail-safe defined in the specification.

Workaround: NA

Status: No Fix

6. Some inbound PCI Express* transactions that require master abort are logged incorrectly

Problem: For specific inbound transactions on PCI Express* that require a “master abort” response, Intel® 631xESB/632xESB I/O Controller Hub is logging in upstream PCI Express* port that it replied with a master abort when that logging is not required. The specific cases are those in which the inbound transaction is master aborted by the PCI Express* downstream port before the transaction reaches the internal switch, but it is logged as though it reached the internal switch before getting master aborted. These are the specific cases: Loopback addresses. BME (bus master enable) disabled. I/O transactions. Config transactions.

Implication: Although the downstream PCI Express* agent logs that it received a master abort, it doesn’t cause a failure to also log this in Intel® 631xESB/632xESB I/O Controller Hub (upstream port interface).

Workaround: NA

Status: No Fix

7. Default value used to calculate L0s exit latency is incorrect.

Problem: The downstream PCI Express* links can operate in either common-clock mode or in non-common-clock mode. When the first PCI Express* port is in non-common-clock mode and the second PCI Express* port is in common-clock mode; then when the second port exits from L0s to L0 it will always go into recovery.

Implication: Takes longer to get into L0

Workaround: BIOS Workaround: Refer to latest Intel® 631xESB/632xESB I/O Controller Hub BIOS spec update.

Status: No Fix

8. Wrong and illegal negotiated width value can be reported for downstream PCI Express* ports

Problem: For the PCI Express* downstream ports (PE1 and PE2), incorrect negotiated link width values can be reported in the following situations:Link goes into recovery and “fails-down” to a smaller width. PCI Express* hot-plug occurs with a different width adapter. Any renegotiation after reset with a different link width. Legal “negotiated width” values for PCI Express* downstream ports are: x8: 01000, x4: 00100, x1: 00001. If a fail-down occurs during recovery, one of the following illegal values can occur: x8 to x4: 01100 (appears as x12), x8 to x4 to x1: 01101 (undefined), x8 to x1: 01001 (undefined), x4 to x1: 00101 (undefined). The other side of the link is not affected and should report the correct width.

Implication: It is a PCI Express* compliance failure to report wrong and undefined negotiated link width values.

Workaround: NA

Status: Fixed

9. HPCCTL register incorrectly located in configuration space

Problem: The HPCCTL (Hot Plug Controller Control Register) is located at offset 90 in the configuration space.

Implication: This register is invisible in normal system configuration.

Workaround: BIOS Workaround: Details in latest Intel® 631xESB/632xESB I/O Controller Hub BIOS spec update.

Status: No Fix

10. Max Payload Size register field has wrong default

Problem: Per the PCI Express* Specification, the “Max Payload Size“of the “Device Control Register” should default to 000b (for 128 bytes max payload size). Intel® 631xESB/632xESB I/O Controller Hub silicon default for both upstream and downstream PCI Express* interfaces is 001b (for 256 bytes). Max Payload Size is a RW value that is set by BIOS after checking the “Supported Max Payload Size” capability (a RO bit) in the “Device Capability Register” for each function in the hierarchy. Legal values for Max Payload Size in Intel® 631xESB/632xESB I/O Controller Hub are “0” (128 bytes) or “1” (256 bytes). There is no problem setting the bit to either 0 or 1 because Intel® 631xESB/632xESB I/O Controller Hub is a switch and only passes through large packets; Intel® 631xESB/632xESB I/O Controller Hub does not generate them. Affects ports PE1 through PE4.

Implication: This is a PCI Express* Specification violation. The PCI-SIG provided PCI Express* Compliance Test will report failure.

Workaround: NA

Status: Fixed

11. Cache Line Size register should be type RW

Problem: Per the PCI Specifications, “Cache Line Size (CLS) Register” should default to type “RW” but in the IOxAPIC configuration space it is type “RO”. The CLS register is not used in IOxAPIC, but for compatibility with legacy PCI the register is defined to be RW. The CLS register for all the other functions in Intel® 631xESB/632xESB I/O Controller Hub is correctly set to RW.

Implication: This is a PCI Specification* violation. The PCI-SIG provided Compliance Test will report failure with IOxAPIC enabled. This bug violates the PCI Express* Specification but is not expected to cause any functional failures

Workaround: NA.

Status: No Fix

12. SATA AHCI version register value incorrect

Problem: Index/Data pair scheme was added to AHCI 1.1 specification; Intel® 631xESB/632xESB I/O Controller Hub implements the Index/Data pair scheme but the SATA AHCI version register has a value of 0001000h indicating AHCI spec 1.0.

Implication: BIOS reading SATA AHCI register will not know that Intel® 631xESB/632xESB I/O Controller Hub supports Index-Data pair.

Workaround: BIOS workaround: See latest Intel® 631xESB/632xESB I/O Controller Hub BIOS spec update.

Status: Fixed

13. SATA EB buffer overflow should set ERR.E bit instead of ERR.M bit

Problem: Intel® 631xESB/632xESB I/O Controller Hub erroneously sets PxSERR:[M] (Recovered Communication Error) when the internal elasticity buffer experiences an overflow or if a miss-align is detected after PhyRDY is detected. A miss-align may happen during resume from the Partial or Slumber link PM states and will cause the M-bit to be erroneously set.

Implication: As a result of M-bit setting while resuming, the host will set the Interface Non-Fatal Error Status (INFS) bit since it happens while the interface has no Frame Information Structure being transferred or received. INFS is expected to be recoverable by S/W (mostly ignored) and should not have an impact to subsequent command execution. No restart of the controller is needed in this case.

Workaround: NA

Status: No Fix

14. Output of the JTAG IDCODE register is a constant ‘1’

Problem: When performing a JTAG IDCODE instruction the output of the IDCODE register of other devices that appear before the Intel® 631xESB/632xESB I/O Controller Hub in the chain will be a constant ‘1’ instead of the actual contents of the IDCODE register. Therefore, it is impossible to read the contents of the IDCODE register.

Implication: Cannot read IDCODE correctly.

Workaround: There are four possible workarounds for A0 1) Don’t do a JTAG IDCODE instruction on Intel® 631xESB/632xESB I/O Controller Hub 2) Put Intel® 631xESB/632xESB I/O Controller Hub JTAG in BYPASS mode 3) Physically bypass Intel® 631xESB/632xESB I/O Controller Hub JTAG via board strappings 4) Put Intel® 631xESB/632xESB I/O Controller Hub first in the JTAG chain.

Status: Fixed

15. Upstream Port captures incorrect value in Device Capabilities register when receives Set_slot_pwr Message

Problem: When Intel® 631xESB/632xESB I/O Controller Hub Switch Upstream Port (PE3 or PE4) receives a SET_SLOT_PWR Message, it does not capture the data payload correctly. The data payload provides the Slot Power Limit Scale and Slot Power Limit Value to be stored in Device Capabilities register (Intel® 631xESB/632xESB I/O Controller Hub upstream port, offset 0x48 [27:26], [25:18]). Due to this bug, software that reads Device Capabilities register will receive incorrect values for Slot Power Limit Scale and Value.

Implication: Impact considered low since Intel® 631xESB/632xESB I/O Controller Hub does not support any slot power budgeting in IO. Intel® 631xESB/632xESB I/O Controller Hub is non-compliant in this situation though, since if doesn't support slot power budgeting, it should return all 0's for these values in Device Capabilities but it doesn't

Workaround: NA

Status: Fixed.

16. Incorrect SET_SLOT_PWR Message data payload on Intel® 631xESB/632xESB I/O Controller Hub external downstream ports

Problem: Intel® 631xESB/632xESB I/O Controller Hub external downstream ports (PE1, PE2) send incorrect values for data payload of SET_SLOT_PWR Message. The Slot Capabilities Register (offset 0x58) Slot Power Limit Scale [16:15], and Slot Power Limit Value [14:7] contain the values that should map to data payload [9:8][7:0] respectively. The mapping is incorrect which results in the payload bits [9:0] being 0's. PCI Express* 1.0a spec in Section 6.9 states that switch downstream ports must not transmit a value lower than lowest value specified by electromechanical spec for slot's form factor. In this case, the payload is transmitting a value of 0.

Implication: PCI Express* specification compliance issue. Possible functional issue; compatibility. Possible worst case is that a PCI Express* adapter would exceed power budget but SW would think it is fine.

Workaround: NA

Status: Fixed

17. Reading LAN PCI function BARS can return incorrect values

Problem: The functions LAN0 and LAN1 each have 3 BARS (MEM, FLASH, and IO). There are two cases that result in bad decoding: 1) If memory BAR (BAR0) of LAN0 or LAN1 is not initialized, then any access to any function will be directed to BAR0 of the same function. 2) If memory BARs of both LAN functions are enabled and flash BAR (BAR1) of LAN0 or LAN1 is disabled then any access to any function will be directed to BAR1 of the same function.

Implication: BARS may contain incorrect values.

Workaround: BIOS Workaround: Must fully initialize LAN0 and LAN1 functions.

Status: Fixed

18. BMC acting as Fast Management Link Master, asserting the FLBINTEX signal can cause loss of synchronization between master and slave.

Problem: When the Intel® 631xESB/632xESB I/O Controller Hub FML port is configured to work as an FML master, the slave may stop the master by pulling down the FLBINTEX pin. The FML e-spec allows asserting FLBINTEX (active low) 53nsec after the last SMBCLK falling edge. Intel® 631xESB/632xESB I/O Controller Hub can stand only 49nsec. If the slave device asserts the FLBINTEX after 49nsec, the synchronization between master and slave will be lost.

Implication: This transaction will probably be repeated forever, depending on the application in the slave device. Most known FML slave devices assert FLBINTEX early enough for Intel® 631xESB/632xESB I/O Controller Hub to sample.

Workaround: NA

Status: Fixed

19. Broadcast Packets Transmitted Counter may wrongly count flow control packets

Problem: All flow control packets that are sent after host transmission of broadcast packet till host transmission of non-broadcast packet are wrongly counted by BPTC (Broadcast Packets Transmitted Count) counter.

Implication: BPTC counter may be incorrect.

Workaround: NA

Status: Fixed

20. Good Packets Received Counter may be incorrect due to Link disconnect

Problem: Good Packets Received Counter (GPRC) statistics counters may count a packet interrupted by Link disconnect. This packet should not be counted due to CRC error

Implication: The Good Packets Received Counter (GPRC) may show a slightly higher number than expected upon the rare events of link disconnect. This counter is only used for statistics, the exact number is not significant.

Workaround: NA

Status: No Fix

21. Intel® 631xESB/632xESB I/O Controller Hub LAN - FML hold time violation per specification

Problem: The FML specification defines the hold time for the FML inputs as 0ns from the falling edge of the clock. When Intel® 631xESB/632xESB I/O Controller Hub LAN is in master mode (in which it generates the clock to the slave), hold time of 2ns should be provided in order for the Intel® 631xESB/632xESB I/O Controller Hub to sample the data correctly.

Implication: Since the Intel® 631xESB/632xESB I/O Controller Hub drives the clock, and the slave identifies a falling edge of the clock before it drives the data, hold time of 0ns is not realistic and would not happen in a real system.

Workaround: NA

Status: No Fix

22. Packets with Symbol error may be counted as good packets

Problem: When a symbol error occurs just after start of packet, the packet is rejected, but the GPRC counter counts this packet as good packet and the RXERRC Counter doesn't count it as a bad packet.

Implication: The implication of this bug is that it will cause both counters values to be slightly inaccurate. This bug is not expected to cause any functional failures.

Workaround: NA

Status: No Fix

23. RNBC counter is not correct

Problem: This counter is intended to count the number of packets received when no host memory buffer is available. When in muliple recieve queue mode, the counter may count false events and may not count real events

Implication: The implication of this bug is that the value of RNBC register does not reflect the number of packets received with no available host buffer.

Workaround: NA

Status: Fixed

24. gtx_clk is driven from the crystal in MII 100/10 mode is not synchronized to 125 MHz internal clock.

Problem: Kumeran block communicates with the MAC using MII/ GMII interface. The MII and Kumeran clocks should be balanced as there is no synchronization level. As designed, the external reference clock supports two input frequencies: 25 MHz and 62.5 MHz. When the external reference clock's input frequency is 25 MHz, GTX_CLK, an internal GMII signal, is driven from the reference clock's input and thus not synchronized to the 125 MHz clock that is used.

Implication: The implication of this bug is that only the 62.5 MHz crystal mode can be used with Kumeran until this bug is fixed (i.e. no support for 25Mhz crystal).

Workaround: NA

Status: Fixed

25. ESI link has no de-emphasis

Problem: De-emphasis is when the 1st bit of a series of bits in a serial data stream is driven to a slightly higher level (for a 1) than the succeeding same type bits (or a slightly lower level for a 0). Verified de-emphasis registers setting for ESI link on Intel® 631xESB/ 632xESB I/O Controller Hub; they are all set to 0 which is disabled (de-emphasis off). De-emphasis should be set on by default when in AC coupled full-swing mode.

Implication: Signal integrity could be degraded.

Workaround: BIOS Workaround: See latest Intel® 631xESB/632xESB I/O Controller Hub BIOS spec update.

Status: Fixed

26. SATA Default Tx Drive Strength set too High at Gen1/Gen2

Problem: Drive strength default settings for Intel® 631xESB/632xESB I/O Controller Hub are too high; causing Vdiffp-p to violate specification.

Implication: Spec violation.

Workaround: BIOS Workaround: See latest Intel® 631xESB/632xESB I/O Controller Hub BIOS spec update.

Status: No Fix

27. Test port pins that are used for straps are not sampled correctly at PE RESET

Problem: When Intel® 631xESB/632xESB I/O Controller Hub LAN is in ARC JTAG state, i.e. the ARC debugger is connected through the JTAG interface, the test port bits 2, 3, 6, and 7are constantly driving out “0”. Upon assertion of PERST# these pins become inputs and should sample the test port strap value which should be “1”. The problem is that the pull-up resistor changes the pin value slowly resulting in a “0” value sampled, and causing glitches on the internal BMC clock.

Implication: The BMC behavior is unpredictable.

Workaround: Board workaround - add tri-state buffer that will strongly drive a “1” when PERST# is asserted.

Status: Fixed

28. PATA config space disappears when secondary IO decode is turned on in native mode

Problem: When the following values are programmed into the PATA (Bus 0h Dev 1Fh Function 1h), the first eight bytes (offset 00h to 07h) of the PCI Configuration space disappear and show the value of 7Fh. Write 07h to offset 04h, write 80h to offset 41h, write 80h to offset 43h, write 8Fh to offset 09h. It seems that the secondary IO decode interferes with the configuration space on config read. The config write still seems to decode correctly.

Implication: Possible problems with IDE driver, causing an inability to boot the IDE drive.

Workaround: BIOS Workaround: See latest Intel® 631xESB/632xESB I/O Controller Hub BIOS spec update.

Status: No Fix

29. PCI Express* lanes not driving to correct level

Problem: PCI Express* downstream ports not driving to correct level. ~900mVp-p on lane 0. All other lanes showing about 300mVp-p. PE1 and PE2 configured as x8:Lane 0 = normal amplitude (~900mVp-p); Lanes 1-7 = low amplitude (~300mVp-p). PE1 and PE2 configured as separated x4s:Lane 0 of PE1 and Lane 0 of PE2 = normal amplitude (~900mVp-p); Lanes 1-3 of PE1 and Lanes 1-3 of PE2 = low amplitude (~300mVp-p).

Implication: This bug violates the PCI Express* Specification but is not expected to cause any functional failures.

Workaround: NA

Status: Fixed

30. IDE setting SRST (software reset) doesn't set the BSY bit.

Problem: Setting SRST doesn't set the BSY bit. According to the IDE spec, when Host asserts IDE SRST command, the BSY bits in IDE Status registers should be set but they are not.

Implication: The IDE driver in windows may disable the IDE function.

Workaround: NA

Status: Fixed

31. MAC Transmission of Preamble violates 802.3 specification.

Problem: In 100FD/HD and 100FD/HD, the DUT transmits only 6 bytes of Preamble instead of 7 bytes. This behavior violates 802.3 spec.

Implication: This is a compliance issue, should not have any affect on network behavior. These packets are received correctly by the link partner.

Workaround: NA

Status: Fixed

32. SATA Gen2 Driver current lower than expected.

Problem: Serial ATA Gen2 main driver current measurements are lower than expected.

Implication: Possible eye violations.

Workaround: BIOS workaround - see BIOS spec update.

Status: No Fix

33. PCI Express* Link Control Register requires word write to update.

Problem: A one byte configuration write at Bm:D0:F3 offset 57h to set slot clock configuration bit fails to update. A word (two bytes) write to set slot clock configuration bit works fine. Affects ports PE1 though PE4.

Implication: Both offsets 56h and 57h must be written to at the same time for a write to take place.

Workaround: BIOS Workaround: See latest Intel® 631xESB/632xESB I/O Controller Hub BIOS spec update.

Status: No Fix

34. PCI Express* downstream ports flag inbound 4K memory read as malformed Transaction Layer Packet (TLP).

Problem: PCI Express* downstream ports (PE1 and PE2) flag inbound 4K memory read requests as malformed TLP. The downstream ports will log the malformed TLP error and discard the memory read requests. This condition could occur if downstream devices can generate read requests up to 4KB.

Implication: Inbound 4K memory read request will fail and the downstream device will lose flow control credit.

Workaround: BIOS Workaround: See latest Intel® 631xESB/632xESB I/O Controller Hub BIOS spec update.

Status: Fixed.

35. PCI Express* Hot Plug: HPCCTL register not sticky.

Problem: The HPCCTL register (Hot Plug Controller Control Register at offset 0x90 in the configuration space) is not sticky across power cycles, but the Slot Capabilties register for Hot-Plug Capable is read/write once sticky (RWOS). Affects ports PE1 and PE2.

Implication: This prevents the correct PCA955x device type from being selected on a warm reset.

Workaround: NA.

Status: Fixed

36. PCI Express* Hot Plug: HPCCTL Start-of-day Power Fault not latched for port B in two single-byte mode.

Problem: When configured for two (dual) single-byte PCA9554 (or PCA9554A) mode, the Intel® 631xESB/632xESB I/O Controller Hub does not latch a start-of-day power fault on port B. Affects ports PE1 and PE2.

Implication: Software would not be aware of a potential start-of-day power-fault and thus would not be able to notify the user why the device in the slot was not working.

Workaround: NA.

Status: Fixed

37. Intel® 631xESB/632xESB I/O Controller Hub unable to enter the PCI Express* Power Management Link States L0s, L1 and L2 on ports PE1, PE2, and Intel® 631xESB/632xESB I/O Controller Hub LAN.

Problem: Intel® 631xESB/632xESB I/O Controller Hub LAN device and PCI Express* downstream ports PE1, PE2 are not able to transition into the L0s, L1 or L2 power management link states. The Intel® 631xESB/632xESB I/O Controller Hub does not issue an acknowledge message to devices on PE1, PE2, or Intel® 631xESB/632xESB I/O Controller Hub LAN that are requesting to go into any of the power management link states.

Implication: 1) ASPM L0s enable register does not enable transmitters to transition into the L0s state on PE1, PE2 PCI Express* links, and Intel® 631xESB/632xESB I/O Controller Hub LAN . 2) PE1, PE2, and Intel® 631xESB/632xESB I/O Controller Hub LAN links may not enter the L1 and L2 states resulting in a system hang. 3) The system can lock up when attempting to go into S3, S4, or S5. 4) System may hang when the operating system or device driver transitions PCI Express* endpoints attached to PE1, PE2, or Intel® 631xESB/632xESB I/O Controller Hub LAN ports into the D3 Hot state. 5)Interchanging pre-installed Windows* OS images between Greencreek LE and Blackford LE based systems may cause a system hang due to device re-discovery and installation which may insert system devices into the D3 Hot state.

Workaround: Partial BIOS workaround for S5. See latest Intel® 631xESB/632xESB I/O Controller Hub BIOS spec update.

Status: Fixed

38. PCI Express* Hot-Plug: Hot swap to lower link width fails.

Problem: Hot-swapping to a lower link width fails on the PCI Express* ports PE1 and PE2. Failure only occurs for changes to lower link widths, changes to higher widths do not fail.

Implication: Will not be able to hot-swap to lower link width.

Workaround: NA.

Status: Fixed.

39. PCI Express* port PE0 Link layer should drop Data Link Layer Packets (DLLPs) with unknown encoding type.

Problem: A received DLLP which is not corrupt, but which uses unsupported DLLP Type encodings is discarded without further action. This is not considered an error.

Implication: If the Intel® 631xESB/632xESB I/O Controller Hub interprets this as a NAK, a needless replay may occur. This bug violates the PCI Express* Specification but is not expected to cause any functional failures.

Workaround: NA.

Status: No Fix

40. Intel® 631xESB/632xESB I/O Controller Hub does not ignore a PCI Express* Null Packet on port PE0.

Problem: If the Intel® 631xESB/632xESB I/O Controller Hub recieves a PCI Express* Null Packet, it should drop the packet and not perform sequence number checking or respond with any ACK or NAK DLLP. The Intel® 631xESB/632xESB I/O Controller Hub still performs sequence number checks for Null packets and may respond with an ACK or NAK depending on the result of the check.

Implication: Intel® 631xESB/632xESB I/O Controller Hub may send ACK or NAK DLLP in response to a Null packet. This may degrade link performance due to unnecessary retries.

Workaround: NA.

Status: No Fix

41. Intel® 631xESB/632xESB I/O Controller Hub sending less than the minimum number of Power Management Acknowledgments (PMAKs) to SATA.

Problem: The SATA spec requires the Intel® 631xESB/632xESB I/O Controller Hub to send at least four PMAKs to the SATA device. Intel® 631xESB/632xESB I/O Controller Hub sends only three PMACKs before entering a lower power state after the device requests partial or slumber state on the SATA bus.

Implication: This bug violates the SATA Specification but is not expected to cause any functional failures.

Workaround: NA.

Status: No Fix

42. Advanced Host Controller Interface (AHCI): Improper length register device-to-host FIS.

Problem: If a SATA device sends less than a five dword Register Device-to-Host FIS, the Intel® 631xESB/632xESB I/O Controller Hub will correctly respond with RERR but may not be able to accept any further FISes from the device.

Implication: SATA bus will hang. This only applies while operating in AHCI mode. No known devices use Register Device-to-Host FIS sizes less than five dwords, as these are not allowed by the Serial ATA Specification, rev 1.0a.

Workaround: The AHCI driver should reset the bus by sending COMRESET, per the AHCI spec.

Status: No Fix

43. Standard Hot-Plug Controller (SHPC) Message Signaled Interrupt (MSI) is lost/corrupted after PCI config access to 0x78 or 0x7C.

Problem: Expected SHPC MSI is not generated for the next hot-plug event after a PCI config access to offset 0x78 or 0x7C (SHPC status and index window to working registers) in the PCI/PCI-X bridge.

Implication: SHPC events would not be detected by the interrupt handler and would not get serviced.

Workaround: NA

Status: Fixed

44. PCI Express* Extended Tag Capability Bit.

Problem: Intel® 631xESB/632xESB I/O Controller Hub incorrectly has the PCI Express* Extended Tags Supported capability bit (D28:F0/1/2/3:Offset 44h:bit 5) set to '1', though the Intel® 631xESB/632xESB I/O Controller Hub does not support Extended Tags. Affects port PE0.

Implication: Software will not be able to implement Extended Tags support.

Workaround: NA.

Status: No Fix

45. Command register Bus Number field attribute not PCI compliant for the PCI-X bridge.

Problem: Register Offset DCh: PCI-X Bridge Status - Bus Number, bits 15:8. This issue was observed using the PCI-SIG Configuration Compliance Test Suite. The Bus Number field should be Read Only for PCI compliance (PCI-X 2.0a spec, section 8.6.2.4), but is Read\Write.

Implication: Register can be overwritten.

Workaround: NA.

Status: No Fix

46. PxSERR errors are logged during BIOS initialization when running at 3Gbps on a small percentage of parts

Problem: This failure signature has only been seen on a few Intel® 631xESB/632xESB I/O Controller Hub parts and only happens when devices are running at 3Gbps. During BIOS boot up as the SATA controller is initialized, PxSERR errors are logged. Error bits 1, 10, and 19 are set which are recovered data integrity error, protocol error, and 10b to 8b decode error. These bits are set in between BIOS postcode 423c and 503c. Some of these failing parts on certain platforms will hang at postcode 0x75 . Affects SATA ports.

Implication: May hang some operating systems\plaforms during boot up.

Workaround: NA.

Status: Fixed

47. Intel® 631xESB/632xESB I/O Controller Hub LAN: When the IDE is diabled by firmware all the BARs return 7F.

Problem: When IDE function is disabled by BMC firmware, BAR4 will also return 7f for all registers in this BAR. The interrupt register in this BAR falsely returns 7f meaning that all bits in the interrupt register are set. When interrupt sharing mode is used and an interrupt occurs, the operating system searches for the device that asserted the interrupt. Due to this bug, Intel® 631xESB/632xESB I/O Controller Hub LAN would appear as a device that asserted the interrupt when it did not.

Implication: May hang some operating systems since the real source of the interrupt would not be found..

Workaround: Leave IDE function enabled.

Status: Fixed

48. PCI Express* Link Training Status and State Machine (LTSSM) on PE1 and PE2 incorrectly sends TS1-PAD-PAD in Configuration.Linkwidth.Start.

Problem: PCI Express* ports PE1 and PE2 LTSSM Transition from Polling.Configuration to Configuration.Linkwidth.Start sends first TS1 ordered set with PAD as link number and should non-PAD. This doesn't match PCI Express* 1.0a spec. When LTSSM is in Polling.Configuration, it sends TS2 ordered set with Link and Lane numbers set to PAD. When the affected port sees 8 consecutive TS2 ordered sets with Link and Lane numbers set to PAD, it goes to Configuration Linkwidth.Start. In this state, the port should send TS1 ordered sets with selected Link numbers and lane numbers set to PAD.

Implication: Bug causes no functional impact with compliant devices.

Workaround: NA.

Status: No Fix

49. Non-Hot Plug Presence Detect bit not set on PCI Express* ports PE1 and PE2.

Problem: For Non Hot-Plug mode, presence detect state bit always reads 0x0, with or without a card in slot and with or without slot implemented bit set. This is true for PCI Express* ports PE1, PE2 and Intel® 631xESB/632xESB I/O Controller Hub LAN.

Implication: SW will always read bit as 0x0.

Workaround: NA.

Status: Fixed.

50. Intel® 631xESB/632xESB I/O Controller Hub LAN not handling Lock Transactions as Unsupported Requests per PCI Express* spec.

Problem: Sending Memory Read Lock Transactions to Intel® 631xESB/632xESB I/O Controller Hub LAN devices results in an Unsupported Request with incorrect Transaction ID and incorrect Completion type. Intel® 631xESB/632xESB I/O Controller Hub LAN port

does not support Lock Transactions which is allowed by PCI Express* 1.0a spec. The issue of incorrect Transaction ID (more specifically the wrong tag field value) will result in the initiator being unable to match the request with the completion. This normally would result in completion timeout & unexpected completion, but the completion timeout must be disabled in any MCH root port which connects to a PCI/PCIX bridge.

Implication: Current LAN / BMC drivers do not send lock transactions. Lock transactions should not be used with Intel® 631xESB/632xESB I/O Controller Hub LAN device.

Workaround: NA.

Status: No Fix

51. Downstream PCI Express* bridges are inaccessible if secondary bus number misconfigured.

Problem: There are a few scenarios which RTL logic results in odd (possibly non-compliant) behavior from a software perspective. These occur when the bus number registers (offset 0x18-Primary, 19-Secondary, 0x1A-Subordinate) are misconfigured or if programmed in a certain order, i.e. non-dword writes.

Implication: SW might be unable to access downstream PCI Express* bridges.

Workaround: See Intel® 631xESB/632xESB I/O Controller Hub BIOS spec update.

Status: No Fix

52. Advanced Host Controller Interface (AHCI) Host Bus Adapter (HBA) BSY bit set when recovering from fatal error.

Problem: During an AHCI fatal error condition, if device signals a Task File Error (TFES), the Intel® 631xESB/632xESB I/O Controller Hub may not be able to recover correctly after software performs the AHCI spec-defined fatal error recovery mechanism.

Implication: SATA port will appear busy resulting in device being inaccessible.

Workaround: AHCI driver should toggle the ST bit to ‘1’ and back to ‘0’ upon detecting TFES bit set after ST bit is cleared.

Status: No Fix

53. Noise on PCI Express* TX coming out of Electrical Idle

Problem: The PCI Express* Common Mode Voltage is not stable immediately after Receiver Detect Sequence when entering Polling.Active from Detect.Active states. Affects ports PE1 through PE4.

Implication: Common Mode Voltage noise may result in bit errors early in Polling.Active state. May result in additional training time before transitioning on to Polling.Configuration.

Workaround: NA

Status: No Fix

54. IDE VIL not meeting spec

Problem: The VIL spec is not being met for IDE signals under some conditions. The VIL spec shown in the EDS is 1V.

Implication: Violation is slightly under spec and not expected to cause any failures.

Workaround: NA

Status: No Fix - Spec update.

55. JTAG read data corrupted when concurrent config access

Problem: For JTAG read cycle to north or southbound PCI Express* ports, the read return data will be placed in the JTAG buffer waiting to be streamed out. If there is a read/write cycle access to any of these ports adress space from PCI Express*, PCIX or SMB before the JTAG data is being streamed out, the buffer will be overwritten with new request data. Affects ports PE1 through PE4.

Implication: JTAG reads may return corrupted data.

Workaround: Do not use JTAG when there is a configuration space access to any of these interfaces or do not access any of the interfaces when a JTAG read is in progress.

Status: No Fix

56. Standard Hot-Plug Controller (SHPC) Hot-Plug PM_PME Requestor ID not correct.

Problem: When the Intel® 631xESB/632xESB I/O Controller Hub is put in D3hot and the SHPC PME enable bit is set, a SHPC event configured to generate an interrupt causes one PM_PME PCI Express* message to be sent before the interrupt. This PM_PME message contains the Requestor ID function number 0 instead of 3.

Implication: Depending on system software, could fail to return from a sleep state.

Workaround: NA.

Status: Fixed

57. PCI Express* Completion timer not halting in L1

Problem: The Intel® 631xESB/632xESB I/O Controller Hub PCI Express* completion timer does not halt when the link enters L1 state. Affects port PE0.

Implication: Intel® 631xESB/632xESB I/O Controller Hub will flag a completion timer error.

Workaround: NA

Status: No Fix.

58. Split-Lock cycle to LPC space resulting in FSB timeout and IERR

Problem: Intel® 631xESB/632xESB I/O Controller Hub may not properly handle split locked cycles to LPC when a PHOLD sequence is going on concurrently. The hang occurs when the second split lock memory read request is sent out on ESI, but never receives a completion.

Implication: System could hang. This issue has only been replicated using a synthetic test environment and has not been reported using commercially available hardware/ software.

Workaround: NA

Status: No Fix.

59. PCI Express* SKP/InitFCx Contention

Problem: During PCI Express* initialization, if a SKP is being transmitted immediately before a InitFCx DLLP, then a partial InitFCx may be transmitted. Affects port PE0.

Implication: A slight delay (less than 100ns) may occur during link initialization. Device may report correctable errors. InitFCx will automatically be repeated.

Workaround: NA

Status: No Fix

60. Upstream PCI Express* Ports (PE3 & PE4) enter the L1 link state when all downstream devices are NOT in the d3hot power management state.

Problem: The upstream PCI Express port transitions into L1 even when some downstream functions/endpoints are in d0 state. Using an Agilent N4220B PCI-E probe monitoring the upstream link to the MCH, it was observed that the link transitions into the L1 state whenever the upstream PCI Express* Port (Bm,D0,F0) and PCI Express*-to-PCIX Bridge (Bm,D0,F3) are programmed to the d3hot state. The IOxAPIC device (Bm,D0,F1) is in the d0 (normal operation) state as are both port PE1 and PE2. No other combination seems to force L1 transition.

Implication: Not expected to cause any functional issues in the system.

Workaround: BIOS workaround. See latest Intel® 631xESB/632xESB I/O Controller Hub BIOS spec update.

Status: No Fix

61. Interrupt Disable attribute does not correspond to implementation for upstream PCI Express* bridge.

Problem: Register offset 04h: Command and Status - Interrupt Disable/Mask, bit 10. This issue was observed using the PCI-SIG Configuration Compliance Test Suite. The Interrupt Disable bit for the upstream PCI Express* switch is implemented as Read-Only with a value of '0'. Section 7.5.1.1 of the PCI Express* 1.0a base spec indicates the attribute should be RW. The most recent PCI Bridge Spec. (June 2003, 1.2) states the bit is optional. Intel feels the PCI Express* spec should explicitly include for virtual bridges the p2p bridge spec's allowance to not implement the Interrupt Disable bit. Affects ports PE3 and PE4.

Implication: NA.

Workaround: NA.

Status: Non Product Issue

62. Max Read Request Size config default not PCI compliant for all PCI Express* ports.

Problem: Register offset 4Ch: PCI Express* Device Control - Max Read Request Size, bits 14:12. This issue was observed using the PCI-SIG Configuration Compliance Test Suite. The Max Read Request Size should have a default value of "010" (512 Bytes) for PCI compliance (section 7.8.4 of PCI Express* 1.0a base spec). Intel® 631xESB/632xESB I/O Controller Hub has a default of "101" (4K bytes). The default is the same for the upstream PCI Express* switch and all three downstream PCI Express* bridges. This default value is interpreted as only for Devices which generate Read requests. Because this part of the Intel® 631xESB/632xESB I/O Controller Hub is a switch and won't generate read requests by itself and also PCI Express* 1.0a spec states that Switches need to pass the packets as is without modifying any fields, it is set to read request default value of maximum of 4KB. Affects ports PE1 through PE4.

Implication: NA

Workaround: BIOS could set the appropriate value.

Status: Non Product Issue

63. Under certain conditions, inbound prefetched PCI read requests may return wrong data to the requestor.

Problem: With some prefetch policy settings, the Intel® 631xESB/632xESB I/O Controller Hub may over-aggressively prefetch data for PCI reads and subsequently return the wrong data to the requestor. This problem only exists when there is more than one active agent on the PCI bus. This problem exists for all supported frequencies. This problem exists on both PCI segments. This problem does not affect PCI-X operation at any supported frequency.

Implication: Inbound read requests that are enabled for prefetching may return invalid data when multiple agents exist on the same PCI bus. No error is reported by the Intel® 631xESB/632xESB I/O Controller Hub.

Workaround: The problem may be corrected via a SW workaround. The workaround corrects the problem at all supported frequencies and all prefetch policy settings. Please refer to the latest version of the *Intel® 6700PXH 64-bit PCI Hub and Intel® 6702PXH 64-bit PCI Hub BIOS Specification* for details on the workaround.

Status: No Fix

64. USB 2.0 Transmit Ring-Back failure.

Problem: USB 2.0 transmit ring-back failures may occur.

Implication: The signal may violate the lower and/or upper part of the eye due to Tx ringback failures

Workaround: See Intel® 631xESB/632xESB I/O Controller Hub BIOS spec update.

Status: No Fix

65. Intel® 631xESB/632xESB I/O Controller Hub LAN: Cannot read from EEPROM by parallel access on some EEPROMS.

Problem: The device cannot read the EEPROM when all the following conditions are met: 1. Parallel access (EERD, 0x00014); 2.Wrong signature (e.g. blank EEPROM); 3.EEPROM size is 512Byte and smaller (8 bit address access).

Implication: EEPROM reads may fail under specific conditions.

Workaround: EEPROM with a wrong signature can be accessed by the bit bang access mode.

Status: No Fix

66. PCI-X clock falling edge slew rate fails spec into test load

Problem: Simulations show that for PCI-X, the falling edge clock slew rate is below the spec value of 1.5 v/ns when run into the spec test load.

Implication: This bug violates the PCI-X Specification but is not expected to cause any functional failures.

Workaround: NA

Status: No Fix

67. Flow Control Completion Data credits may incorrectly update based on length field

Problem: Inbound transactions that result in an Unsupported Request response should have the length field zeroed out. A non-compliant device may not zero out this field. The Intel® 631xESB/632xESB I/O Controller Hub should ignore the value in the length field but does not. This results in Intel® 631xESB/632xESB I/O Controller Hub ports PE1 through PE4 Flow Control Completion tracking issues because to Intel® 631xESB/ 632xESB I/O Controller Hub uses length field on CPL packets for tracking of data credits.

Implication: A non-compliant device returning an Unsupported Request response with the length field not zeroed out could cause the Flow Control Completion data credits to be incorrectly updated.

Workaround: NA

Status: No Fix

68. Intel® 631xESB/632xESB I/O Controller Hub LAN: Missed packets in 10/100Mb HD

Problem: When the device operates in 10/100Mb HD with multiple-requests or Large Send enabled, there could be received packet loss. When the Tx FIFO is full, the Tx flow may block the host DMA interface of the device. When the transmission of packets is prevented for a long time, due to capture effect or very long backoff, the transmit FIFO is filled and the fetch of Rx descriptors may be prevented. This will prevent the release of the packets from the Rx FIFO to the host, causing the Rx buffer to overflow and the loss of incoming packets. This is a temporary state that will be released once the transmit side is able to empty the Tx packet buffer.

Implication: There could be some packet loss in the Rx path. Normally, these packets will be retransmitted by upper layer protocols.

Workaround: NA.

Status: No Fix

69. Poisoned TLP generated on PCI Express* MSI hot plug interrupt

Problem: Message is a kernel panic and a reported data parity error on the FSB; kernel panic after pressing attention button for hot plug on slot 2 when slot power is off. Problem exists on PE1 and PE2.

Implication: May get a data parity error reported on the front side bus.

Workaround: Use the legacy interrupts or GPE Message feature for all PCI-Express* Hot-Plug events.

Status: No Fix

70. Intel® 631xESB/632xESB I/O Controller Hub LAN: Certain values of last ip_option may corrupt RSS decision

Problem: When IPV6 tunneled-in-IPV4 packets are received, IP options with data are present, and the last byte of the IP options is 0x08, the value of the RSS hash may be incorrectly set to ‘0’ and the queue and CPU numbers may be incorrectly calculated.

Implication: When working with RSS, the platform uses the RSS hash to do TCP context lookup and has no way of recovering if the RSS hash value is incorrect. In this case, it will drop the packet, and may possibly reset the connect. In addition, this packet may be directed to a wrong queue and wrong CPU.

Workaround: Software workaround: If RSS hash value is 0 and PKTTYPE = 3,4,9, or 0xA, check IP length. If options are present, do not indicate an RSS hash value to the stack. The TCP stack will calculate the RSS hash value for a TCP packet, which will prevent it from being dropped. This workaround is implemented in all Intel LAN drivers.

Status: No Fix

71. PCI Express* Hot-Plug interrupt not sent simultaneously with PM_PME

Problem: In response to a PCI Express* Hot Plug event, the Intel® 631xESB/632xESB I/O Controller Hub does not send a PCI Express* hot-plug interrupt simultaneously with the PM_PME message, but may delay until the software (PME handler) clears the PME Status register bit.

Implication: Software model should not depend on the interrupt being sent along with PM_PME messsage. However, if the PME driver is capable of clearing the PME_STS bit then the hot-plug interrupt can be expected after the PME Status Register bit is cleared.

Workaround: NA

Status: No Fix

72. SATA protocol errors detected in PxSERR register during boot

Problem: Intermittent SATA protocol errors are being observed in PxSERR register on boot. The error can be cleared and may not happen after a reboot. The error condition has only been observed when using Gen2 speed and affects all 6 SATA ports.

Implication: Intel® 631xESB/632xESB I/O Controller Hub may log PxSERR during boot.

Workaround: See Intel® 631xESB/632xESB I/O Controller Hub BIOS spec update.

Status: No Fix

73. Intel® 631xESB/632xESB I/O Controller Hub LAN: Invalid commas on the Serdes interface corrupts the data integrity of the preceding packet

Problem: When receiving valid packet followed by commas with invalid bits, the packet may not be counted in the Good Packets Received Counter. The packet reception is not affected.

Implication: The Good Packets Received Counter may show a slightly lower number than expected upon the rare event of invalid commas reception. This counter is only used for statistics, the exact number is not significant.

Workaround: NA.

Status: No Fix

74. SATA BSY timeout after Standby Immediate and SRST (after 1ms)

Problem: Devices with long slumber latency greater than 10ms may not be available returning from Standby. The BSY bit remains set in IDE interface (i.e. device unavailable) after SRST sent. However, the BSY bit will be cleared if SRST sent twice.

Implication: The device may not appear available to software until SRST is retried.

Workaround: See Intel® 631xESB/632xESB I/O Controller Hub BIOS Spec Update.

Status: No Fix

75. Unsolicited COMINIT while FIS posting pending will corrupt the FIS posting cycle

Problem: In only AHCI mode, the SATA controller may post a FIS incorrectly if either an unsolicited COMINIT arrives or if software performs a port reset, while FIS posting is pending internally.

Implication: Both a malformed TLP is delivered and inappropriate data is delivered on the next upstream cycle.

Note: This has only been replicated in a synthetic test environment and has not been reproduced in production environment.

Workaround: None. The system can be configured to detect this anomalous condition and reset the system to prevent this data migration. See Intel® 631xESB/632xESB I/O Controller Hub BIOS Spec Update.

Status: No Fix

76. Link Down at Upstream PCI Express* is resulting in failure of training on the Downstream PCI Express*

Problem: A link down at Intel® 631xESB/632xESB I/O Controller Hub upstream PCI Express*, either by Secondary Bus Reset (SBR), Link Disable, or forced LTSSM change, may result in failure of training on the downstream PCI Express* links. Once the failure occurs, the downstream links can not be restored without a fundamental reset.

Implication: Software should not perform a secondary bus reset to Intel® 631xESB/632xESB I/O Controller Hub upstream PCI Express* port.

Workaround: NA.

Status: No Fix

77. Link Control Retrain Link bit not reading 0 when slot is empty

Problem: The PCI Express* 1.0a Link Control Register Retrain Bit, bit [5], is a RW, but a read of this bit should always return 0b. Intel® 631xESB/632xESB I/O Controller Hub PCI Express* down-stream ports PE1 and PE2 only function correctly when a card is in the slot, i.e. When the slot is empty, a write of 1b to this bit followed by a read will return 1b instead of 0b as required by the PCI Express* 1.0a spec.

Implication: Incorrect Retrain Link bit returns value of 0b when slot is empty, but no impact to retrain link bit functionality.

Workaround: NA.

Status: No Fix

78. Intel® 631xESB/632xESB I/O Controller Hub LAN: Return loss in SERDES

Problem: The SERDES transmitter's differential return loss is up to -10dB instead of -15dB requirement (PCI Express* sets -10dB as the requirement instead of the previous -15dB in the base 1.0a specification. PICMG3.1 Specification is based on that requirement).

The SERDES receiver's differential return loss is up to -14dB instead of -15dB requirement.

Implication: NA.

Workaround: NA.

Status: No Fix

79. A PCI Express* Downstream ports detect correctable errors while running L0s

Problem: Correctable errors may occur when L0s is enabled on PCI Express* downstream ports or the corresponding device below these links. The correctable errors that could be logged are Receiver errors, Bad TLP, Bad DLLP, Replay Timeout and occasionally Replay Rollover.

Implication: Correctable error may occur on PCI Express* downstream ports when L0s is enabled.

Workaround: Do not enable L0s on the PCI Express* downstream ports or the corresponding device below these links.

Status: No Fix

80. PCI-Express* Hot-Plug slot disable when downstream link is in L0s or L1 can hang the system

Problem: A PCI Express* port receiver can not distinguish an “electrical idle” state that is due to a low power link state versus a link down state. A transaction targeting a link in this situation could result in a dropped packet per PCI Express* 1.0a specification regarding DL_Down and result in potential system hang. The PCI-SIG specification does not clearly cover the methodology for an implementation to handle hot-removal of an end device while the link is in a low power management link state (e.g. L0s or L1).

Implication: Potential unrecoverable system hangs when a device is hot-removed from an Intel® 631xESB/632xESB I/O Controller Hub PCI Express* downstream port while the link is in a low power management link state (i.e. L0s or L1).

Workaround: There are two workaround options:

- Do not enable L0s or L1 (D3hot) on the PCI Express* downstream ports when hot-plug slot disable operations are expected; or
- Use the “Link Disable” functionality of the “Link Control” register in the PCI Express* downstream ports (Bp:D0/1:F0) as part of the process for PCI Express* hot-plug slot enable/disable operations (i.e. slot power control). Specifically, set bit 4 “Link Disable” of Bp:D0/1:F0 offset 0x54 “Link Control” register before setting bit 10 “Power Controller Control” of Bp:D0/1:F0 offset 0x5C “Slot Control” register; and conversely, clear the “Link Disable” bit after clearing the “Power Controller Control” bit.

Status: No Fix

81. Intel® 631xESB/632xESB I/O Controller Hub LAN: Link fail in middle of Rx packet might cause data corruption

Problem: When operating in 10Mb/s and there is a link fail in the middle of a receive packet, the first packet after the link is up might be concatenated with the previous fragment.

Implication: Data may be corrupted in the packet received after recovering from the link loss event.

Workaround: When operating in 10Mb/s, the MAC should be reset upon link fail, in order to clean the internal memory from any fragment that might have been a result of the link fail event.

Status: No Fix

82. Intel® 631xESB/632xESB I/O Controller Hub LAN - Does not resent pause packet

Problem: When flow control is enabled, and Intel® 631xESB/632xESB I/O Controller Hub LAN is congested, it will send a pause packet. If the pause packet timeout expires Intel® 631xESB/632xESB I/O Controller Hub LAN will not send a new pause packet even if it is still congested. If a new packet is received in this condition, this packet will be dropped and a new pause packet will be sent to the link partner.

Implication: In severe and long congestion state, some packets will still be lost when flow control is enabled. The packet loss will be in very low percentage. Packets that were lost will be retried by upper layer protocols.

Workaround: None. Setting the TTV register to its max value will minimize the effect.

Status: No Fix

83. Intel® 631xESB/632xESB I/O Controller Hub LAN - Ethernet frame length issues

Problem: Intel® 631xESB/632xESB I/O Controller Hub LAN does not report of a length error if the 802.3 length field does not match the length of the frame that was received, nor does it discard the packet. In addition, the Intel® 631xESB/632xESB I/O Controller Hub LAN does not strip the padding from frames that have a length field smaller than 0x4E.

Implication: Although the IEEE 802.3 requires the MAC to strip the padding and discard packets with length error, no real system level failure will occur. The Ethernet 2 protocol, on which the 802.3 spec is based did not define any length field. The mode of operation where a length field does not exist is still supported by the 802.3 standard. Thus higher layer protocols always knew and still know how to deal with this issue by using their own length field.

Workaround: NA

Status: No Fix

84. Intel® 631xESB/632xESB I/O Controller Hub LAN - IDE BAR4 offset 2

Problem: The BIOS uses bits 6:5 in offset 2 of BAR4 of the IDE function in order to notify to the operating system whether the IDE controller that is behind this function is capable of performing bus master operation. In Intel® 631xESB/632xESB I/O Controller Hub LAN there is a shift of one bit left when writing to this register, thus a wrong value is read by the operating system.

Implication: If the BIOS enables IDE bus master capabilities, the operating system may not get the correct indication. At this point, the Intel® 631xESB/632xESB I/O Controller Hub LAN FW does not support IDE master capabilities, thus this bug has no meaning. This bug could impact the IDE master operation if a FW that would support this capability will be written in the future. In this case, the operating system will not use the master capabilities, although it would be supported by the FW.

Workaround: BIOS should write a value that is one bit shifted right to this register, e.g. instead of writing 0x20, it should write 0x10. This work around will be required only if a FW that supports IDE master capabilities will be written in the future.

Status: No Fix

85. Intel® 631xESB/632xESB I/O Controller Hub LAN - FML - clock extending on the stop cycle

Problem: The FML protocol defines a mechanism of flow control, in which the slave device can extend the clock by asserting the FMLINTEX pin. The slave uses this mechanism when it can not meet the throughput in which data is transferred. This will temporarily halt

the transaction. When the Intel® 631xESB/632xESB I/O Controller Hub LAN is used as the master device and the slave device asserts the FMLINTEX pin on the stop cycle, the Intel® 631xESB/632xESB I/O Controller Hub LAN ignores it and toggles the clock one more time to complete the transaction.

Implication: The stop cycle occurs after the transaction was already completed and no data is being transferred. Thus there is no much meaning to assert the FMLINTEX at this point and it is not expected to have a slave device asserting the FMLINTEX on the stop cycle.

Workaround: NA

Status: No Fix

86. Intel® 631xESB/632xESB I/O Controller Hub LAN - SERDES is unable to acquire synchronization from ordered sets beginning with /K28.1/ and /K28.7

Problem: According to the IEEE802.3 standard, the device should acquire synch when ordered set begin with K28.1, K28.5 and K28.7; however, the Intel® 631xESB/632xESB I/O Controller Hub LAN acquires synch only when the ordered set starts with K28.5.

Implication: K28.1 and K28.7 symbols are not expected in the network, as the transmitter is only allowed to send K28.5 symbol. K28.5 is transmitted by the link partner for a long time, thus even if a K28.5 symbol is corrupted due to error on the line, and received as K28.1 or K28.7, the device will still acquire synch due to other K28.5 symbols.

Workaround: NA

Status: No Fix

87. Intel® 631xESB/632xESB I/O Controller Hub LAN - Formed and invalid /C/ code handling

Problem: The device may respond improperly to certain invalid sequences, which include comma characters different than K28.5 or symbols with inverted disparity. The device may restart auto-negotiation when it should not, it may not restart auto-negotiation when it should.

Implication: In normal operation the comma used is K28.5 and flipped disparity should not happen on a normal system.

Workaround: NA

Status: No Fix

88. Intel® 631xESB/632xESB I/O Controller Hub LAN - False detection of an idle_match condition

Problem: The idle_match function is used during the auto-negotiation process. This function continuously indicates whether three consecutive /I/ ordered sets have been received and it is observed when moving from IDLE_DETECT state to LINK_OK state within the auto-negotiation state machine. The issue is that even though there are not three consistent /I/ symbols (that is,. there is some combinations of /I/ and other symbols) the device can incorrectly set the idle match to true.

Implication: This failure should not be seen in normal-use cases where there are many consecutive /I/ symbols in the auto-negotiation process. However, if the erroneous case occurs the auto-negotiation will continue and lock on the next /I/ pattern.

Workaround: NA

Status: No Fix

89. Intel® 631xESB/632xESB I/O Controller Hub LAN - Ability Match and Acknowledge Match

Problem: The device does not reset its match count upon reception of /I/ ordered sets in between /C/ ordered sets. Thus the device could reach ability_match, or acknowledge_match state when the /C/ ordered sets that were received were not consecutive, but separated by /I/ ordered sets.

Implication: In normal behavior, the link partner will only send consecutive /C/ ordered sets, thus this scenario will never happen in normal operation.

Workaround: NA

Status: No Fix

90. Intel® 631xESB/632xESB I/O Controller Hub LAN - ROM FW - Bus may hang when Master reads more bytes than slave reported

Problem: With Intel® 631xESB/632xESB I/O Controller Hub, when working in I2C protocol, and when an external BMC executes I2C read transaction, EB2 responds with block of data in which the first returned byte indicates data-length. If BMC attempts to read more bytes than specified by the data-length byte, bus hang may occur.

Implication: External BMC operating in I2C mode that reads slave data disregarding the data-length, will cause bus to hang.

Workaround: When external BMC acts as Master, BMC should interpret the first returned data byte received from Intel® 631xESB/632xESB I/O Controller Hub as data-length, and should stop transaction after reading specified number of bytes.

Status: No Fix

91. Intel® 631xESB/632xESB I/O Controller Hub LAN - ROM FW - SOL Timeout character control byte in EEPROM image does not function

Problem: In Intel® 631xESB/632xESB I/O Controller Hub Gigabit Ethernet Controller, SOL (Serial Over LAN) character control can be configured from the EEPROM. Packets from host to management console will be sent when either the maximum buffer size or a timeout are reached. However, instead of restarting the timer on every transmit to LAN, the timer is restarted every time a new character arrives from the host. When the transmit rate from host is slow, the characters will only be sent when the buffer threshold is reached.

Implication: Characters transmitted from the host may not arrive at the remote console at the expected refresh rate but at bursts. This will usually be noticed only at slow rates (e.g. manual typing), which is not a use case in SOL.

Workaround: NA

Status: No Fix

92. Intel® 631xESB/632xESB I/O Controller Hub LAN - ROM FW - BMC fragments that are sent through 2 different SMBUS ports are sent over LAN as a single packet

Problem: In PT mode, when an external BMC sends sequential fragments of two packets using different SMBUS ports of Intel® 631xESB/632xESB I/O Controller Hub, they are linked and transmitted as one packet.

Implication: External BMC cannot send two unsynchronized packets over two ports.

Workaround: With Intel® 631xESB/632xESB I/O Controller Hub, external BMC should send only one packet at a time.

Status: No Fix

93. Intel® 631xESB/632xESB I/O Controller Hub LAN - Time until Jam is longer than it should be

Problem: When a collision occurs during packet transmission, the MAC is expected to terminate the transmission with a JAM pattern. The Intel® 631xESB/632xESB I/O Controller Hub LAN terminates the transmission with a legal JAM pattern, however, the delay between the collision event and the JAM transmission is longer than defined in 802.3 specification.

This could imply that late collision may occur in a network with an Intel® 631xESB/632xESB I/O Controller Hub LAN running at 10/100 Mb/s half duplex.

Implication: Though this failure implies that late collisions may occur in the network, testing of the Intel® 631xESB/632xESB I/O Controller Hub LAN in maximum Ethernet topology with various hubs did not show any significant increase in late collision count with respect to the same topologies without the Intel® 631xESB/632xESB I/O Controller Hub LAN.

Workaround: None.

Status: No Fix

94. Configuration cycle failure when targeting PCI Express* PE1 and PE2 while PM L1 entry sequence is in process

Problem: PCI Express* switch downstream ports PE1 and PE2 may not completely block TLP and may send a corrupted configuration transaction during the PM L1 protocol during a downstream device's PM L1 re-entry. The PM L1 re-entry behavior, termed "aggressive PM L1", implies devices re-enter PM L1 based on the PMCSR D-state setting. This potential exposure applies only to devices that exhibit the aggressive PM L1 behavior and no failures have been observed with "less aggressive PM L1" devices.

Implication: Configuration cycles targeting an aggressive PM L1 endpoint device during PM L1 re-entry could result in a system hang.

Workaround: See Intel® 631xESB/632xESB I/O Controller Hub BIOS Spec Update.

Note: Depending on specific boot configuration, this workaround may result in a decrease in performance during boot only

Status: See Intel® 631xESB/632xESB I/O Controller Hub BIOS Spec Update

95. PCI Express* ports may gate PM L1 entry based on lack of Flow Control Credit availability

Problem: PCI Express* 1.0a Errata C7, ratified September 25, 2003 strikes through the requirements for upstream posts to wait for minimum credits before proceeding to PM L1. Intel® 631xESB/632xESB I/O Controller Hub PCI Express* ports (PE1 and PE2) are implemented without PCI Express* 1.0a Errata C7 and wait for minimum credits before proceeding to PM L1.

Implication: No impact expected when using PCI Express* 1.0a compliant end-point devices which return flow control credits as described by the PCI Express* specification. PCI Express* end-point devices which are not compliant with the PCI Express* specification or do not return flow control credits after a configuration cycle is received, may cause system instability, degraded performance, or system hangs. Failures are not expected to occur with PCI Express* compliant end-point devices..

Workaround: None.

Status: No Fix

96. ESB2 SATA Signal Voltage Level

Problem: Intel® ESB2 Serial ATA (SATA) transmit buffers have been designed to maximize performance and robustness over a variety of routing scenarios. As a result, the Intel® ESB2 SATA transmit signaling voltage levels may exceed the maximum motherboard TX connector and device RX connector voltage specifications (Section 6.6.2 of SATA Specification, rev 1.0a).

Implication: None known.

Workaround: None.

Status: No Fix

97. Link control register bit 3 in PCI Express Capability table should be read-only and always return 0 for switch ports

Problem: Device 3500 is the upstream port of a PCIe switch. Device implements Read Completion Boundary as R/W even though this bit should be 0 and RO for switch. This bit is not applicable for switches per PCIe Base Spec. 1.0a

Implication: This bit has no functionality for switches per spec. Although device implements this as R/W, it has no logic connectivity beyond the register bit. Considered a non-issue.

Workaround: None.

Status: No Fix

98. Device control register bit 8 should be read-writable if Extended Tag Field Support is indicated

Problem: Device advertises Extended Tag Capability but does not support Extended Tag Field Enable in the Device Control register.

Implication: As a requester this device will only support a 5 bit Tag Field.

Workaround: None.

Status: No Fix

99. CPU IERR during POST and CPU IERR with hang at OS prompt

Problem: Two CPU IERR issues surface (CPU IERR on POST and CPU IERR with hang at OS login prompt), while continuously rebooting system.

Implication: CPU IERR appears due to BIOS not reading the Other Interrupt Control (OIC) register (RCBA + 31FFh) after the APIC enable (AEN) bit is modified.

Workaround: The Intel® 631xESB/632xESB I/O Controller Hub requires the Other Interrupt Control (OIC) Register be read after the APIC Enable (AEN) bit is modified. This read should occur prior to any access to the IOxAPIC address range.

Status: No fix.

100. ESB2 PCIe switch mis-formats TLPs with ECRC destined to the MCH

Problem: In the event that ECRC (End-to-End CRC) is enabled on PCIe endpoints connected to the downstream ports (non-ESI ports) of the ESB2 switch, the ESB2 PCIe switch inserts an extra transaction digest for I/O and configuration write completions destined to the MCH root port. Switches are expected to forward TLPs with ECRC untouched. This issue has only been observed with register attribute tests that do not compare the capabilities of the MCH PCIe root port and endpoints before enabling ECRC. This issue is not observed with PCIe ports of the ESI portion of ESB2. Note that all MCH (Memory Controller Hubs) that are compatible with the ESB2 device do not support ECRC and do not advertise ECRC capability in their Advanced Error Control and Capabilities register.

As a result device drivers, firmware or applications are not expected to enable this capability.

Implication: Malformed TLP status will be asserted by the MCH PCIe root port on receipt of I/O or configuration write completions. Additionally, a system hang may occur if the transaction completion time out is disabled on the MCH root port.

Workaround: Set bit 1 of the upstream PCIe port (Bus m Device 0 Function 0 offset 42h) soon after PCI enumeration. This bit will also need to be restored for workstation implementations in the S3 resume path. Setting of this bit excludes the transaction digest field for all TLPs destined to the MCH root port when ECRC is enabled on the endpoint.

Status: No Fix.

101. Enabling RSS in the Middle of Received Packets May Stop Receive Flow

Problem: The RSS feature can be enabled by the software only during the MAC initialization phase. If, during the time in which RSS_EN is set, unicast packets with the MAC's destination address are being received, the receive path may stop.

Implication: Normally RSS would be enabled when the driver is loaded. At this time, the MAC address is not known to the devices in the network and the only communication with the device would be done through broadcast packets (for example ARP packets would be sent in order to find the MAC address of this new MAC in the network). Thus, no problem will be seen. In the cases when the driver resets the device and re-initializes it when it is already known in the network and receives packets (for example, TCP connections are active), the device may stop on the RSS enable command.

Workaround: The following actions will alleviate this condition:

- Keep manageability setup before change: Read MANC register and keep value of TCO_RCV_EN and RCV_ALL bits
- Direct traffic to manageability: Set MANC.RCV_ALL and TCO_RCV_EN bits before a software reset (use the software/firmware register access semaphore)
- Initiate the software reset.
- Disable host traffic reception: Disable the RAH0/RAL0 valid bit.
- Restore manageability setup as before SW reset: Restore MANC.RCV_ALL and TCO_RCV_EN bits before a software reset (use the software/firmware register access semaphore)
- Enable RSS.
- Configure Rx (address tables, registers etc.)
- Enable host traffic reception: Enable the RAH0/RAL0 valid bit.

Status: No Fix.

5 Specification Changes

There are no new Specification Changes in this revision of the Specification Update.

6 Specification Clarifications

1 PCI Downstream Device Disable Clarification

In Section 5.3.5 PCI-to-PCI Bridge Model, the following note is added:

Note: All downstream devices should be disabled before reconfiguring the PCI Bridge. Failure to do so may cause undefined results.

7 Documentation Changes

1. Content updated to Section 25.1.34

25.1.34 PCI Express_CAP — PCI Express* Capability Register(Bn:F0/F1)

Offset Address: E2–E3h Attribute: RO
Default Value: 0001 Size: 16 bits

Bit	Description
15:14	Reserved.
13:9	**Interrupt Message Number(INT_MSG_NUM)**— RO. Multiple MSI per function is not supported. Hardwired to 00000b
8	**Slot Implemented** — RO. Slot option not implemented hardwired to 0.
7:4	**Device/Port Type** — RO. Indicates the type of PCI Express functions. Hardwired to ~~0001b,~~ 0000b
3:0	**Capability version(CAP_VERSION)** — RO. Indicates the PCI Express capability structure version number 0001b.

2. Content updated to Section 25.2.34

25.2.34 PCI Express_CAP — PCI Express* Capability Register(Bn:F2)

Offset Address: E2–E3h Attribute: RO
Default Value: ~~0011,~~0001/0011h Size: 16 bits

Bit	Description
15:14	Reserved.
13:9	**Interrupt Message Number(INT_MSG_NUM)**— RO. Multiple MSI per function is not supported. Hardwired to 00000b
8	**Slot Implemented** — RO. Slot option not implemented hardwired to 0.
7:4	**Device/Port Type** — RO. ~~Indicates the type of PCI Express functions. The value of for IDE Redirection Function is loaded from the EEPROM word 19h bit 13.~~ Indicates the type of PCI Express functions. The value of the IDE function is loaded from the EEPROM word 19h bit 13 and can be either 0000 b or 0001.
3:0	**Capability version(CAP_VERSION)** — RO. Indicates the PCI Express capability structure version number 0001b.

3. Content updated to Section 25.3.32

25.3.32 PCI Express_CAP — PCI Express* Capability Register(Bn:F3)

Offset Address:	E2–E3h	Attribute:	RO
Default Value:	~~0011,~~0001h	Size:	16 bits

Bit	Description
15:14	Reserved.
13:9	**Interrupt Message Number(INT_MSG_NUM)**— RO. Multiple MSI per function is not supported. Hardwired to 00000b
8	**Slot Implemented** — RO. Slot option not implemented hardwired to 0.
7:4	**Device/Port Type** — RO. Indicates the type of PCI Express functions. Hardwired to ~~0001b,~~ 0000b
3:0	**Capability version(CAP_VERSION)** — RO. Indicates the PCI Express capability structure version number 0001b.

4. Content updated to Section 25.4.33

25.4.33 PCI Express_CAP — PCI Express* Capability Register(Bn:F4)

Offset Address:	E2–E3h	Attribute:	RO
Default Value:	~~0011,~~0001h	Size:	16 bits

Bit	Description
15:14	Reserved.
13:9	**Interrupt Message Number(INT_MSG_NUM)**— RO. Multiple MSI per function is not supported. Hardwired to 00000b
8	**Slot Implemented** — RO. Slot option not implemented hardwired to 0.
7:4	**Device/Port Type** — RO. Indicates the type of PCI Express functions. Hardwired to ~~0001b,~~ 0000b
3:0	**Capability version(CAP_VERSION)** — RO. Indicates the PCI Express capability structure version number 0001b.

5. Content updated to Section 25.5.35

25.5.35 PCI Express_CAP — PCI Express* Capability Register(Bn:F3)

Offset Address:	E2–E3h	Attribute:	RO
Default Value:	~~0011,~~0001h	Size:	16 bits

Bit	Description
15:14	Reserved.
13:9	**Interrupt Message Number(INT_MSG_NUM)**— RO. Multiple MSI per function is not supported. Hardwired to 00000b
8	**Slot Implemented** — RO. Slot option not implemented hardwired to 0.
7:4	**Device/Port Type** — RO. Indicates the type of PCI Express functions. Hardwired to ~~0001b,~~ 0000b
3:0	**Capability version(CAP_VERSION)** — RO. Indicates the PCI Express capability structure version number 0001b.

6. Content updated to Section 25.6.33

25.6.33 PCI Express_CAP — PCI Express* Capability Register(Bn:F4)

Offset Address: E2–E3h
Attribute: RO
Default Value: ~~0011,~~0001h
Size: 16 bits

Bit	Description
15:14	Reserved.
13:9	**Interrupt Message Number(INT_MSG_NUM)**— RO. Multiple MSI per function is not supported. Hardwired to 00000b
8	**Slot Implemented** — RO. Slot option not implemented hardwired to 0.
7:4	**Device/Port Type** — RO. Indicates the type of PCI Express functions. Hardwired to ~~0001b,~~ 0000b
3:0	**Capability version(CAP_VERSION)** — RO. Indicates the PCI Express capability structure version number 0001b.

7. Content updated to Table4-1

PXIRQ[11]#	~~Pull-up~~ None	~~8.33K~~ N/A	13
PXIRQ[12]#	~~Pull-up~~ None	~~8.33K~~ N/A	13

8. Content added to Section 5.25

5.25 Forwarding Cycles to LPC

ESI cycles in the LT register space's FED4_C000h through FEF4_xFFF0 range are to be forwarded to the LPC I/F. These cycles are performed using two new LPC special cycles: LPC TPM-Write and LPC TPM-Read.

Note: The names TPM-Write and TPM-Read are used because external documentation to the TPM vendors calls them this. The exact same cycle type is used for access to the TMKBC.

5.25.1 LPC TPM Write Cycle Format

The new TPM-Write special cycle format is shown in Table 5-80. It is similar to the existing LPC 16-bit I/O write. This allows the LPC LT Register space to be 64 Kbytes.

- If the CPU attempts to write to more than one byte at a time to the LPC Device (TPM or TMKBC), the chipset breaks this up into multiple 8-bit cycles.
- The CPU is permitted to write 1–8 bytes to LPC.

— The ICH may receive either a 4 byte or 8 byte access on ESI.

— The writes may have any combination of BEs set within the 4 or 8 byte access.

— The cycles must be naturally aligned.

Table 5-80. New LPC Special Cycle TPM-Write for Accessing TPM or TMKBC

Field	Value for Bits [03:00]	Description
START	0101	This was a previously reserved value. Now allocated for LT TPM-Write and LT TPM-Read.
CYCTYPE+DIR	0010	Same as used for standard LPC I/O write.
ADDR	See Description	Four nibbles. The LT register space on LPC is 16-bits.
DATA-Low	DIGEST low nibble	
DATA-High	DIGEST high nibble	
TAR		Standard LPC TAR
SYNC		Standard SYNC field for an I/O Write
TAR		Standard LPC TAR

5.25.2 LPC TPM Read Cycle Format

The new TPM-Read special cycle format is shown in Table 5-81. It is similar to the existing LPC
16-bit I/O read. This allows the LPC LT Register space to be 64 Kbytes.

- If the CPU attempts to read more than one byte at a time from the LPC Device (TPM or
TMKBC), the chipset breaks this up into multiple 8-bit cycles.
- The CPU is permitted to read 1–8 bytes from the LPC Device.

— The ICH may receive either a four byte or eight byte access on ESI.

— The read may have any combination of BEs set within the four or eight byte access.

— The cycles must be naturally aligned.

Table 5-81. New LPC Special Cycle TPM-Read for Accessing Fixed Token

Field	Value for Bits [03:00]	Description
START	0101	This was a previously reserved value. Now allocated for LT TPM-Write and LT TPM-Read.
CYCTYPE+DIR	0010	Same as used for standard LPC I/O write.
ADDR	See Description	Same as TPM-Write
TAR		Standard LPC TAR
SYNC	Standard	Standard SYNC field for an I/O Write
DATA-Low	DIGEST low nibble	
DATA-High	DIGEST high nibble	
TAR		Standard LPC TAR

9. Content modified to section 23.1.34

23.1.34 PCS—Port Control and Status Register (SATA–D31:F2)

~~Address Offset: 91h–93h~~ ~~Attribute: R/W, R/WC, RO~~
~~Default Value: 0000h~~ ~~Size: 24 bits~~

~~Bits~~	~~Description~~
~~7:6~~	~~SATA Mode Select: Software programs these bits to control the mode in which the SATA HBA should operate:~~ ~~00 IDE mode~~ ~~01 AHCI Mode~~ ~~10 RAID mode~~ ~~11 Reserved~~ ~~NOTES:~~ ~~1. When combined mode is used (non-zero MV), only IDE mode is allowed. IDE mode can be selected when RAID is enabled.~~ ~~2. AHCI mode may only be selected when MV=0. AHCI mode may be selected when RAID is enabled.~~ ~~3. RAID mode may only be selected when MV=0. Note that RAID5 mode may not be supported in certain SKUs of Intel® 631xESB/632xESB I/O Controller Hub.~~ ~~Programming these bits with values that are illegal (for example, selecting RAID when in combined mode) will result in-deterministic behavior by the hardware.~~
~~5:2~~	~~Reserved.~~
~~1:0~~	~~Map Value — R/W. Map Value (MV): The value in the bits below indicate the address range the SATA ports responds to, and whether or not the PATA and SATA functions are combined. When in combined mode, the AHCI memory space is not available and AHCI may not be used.~~ ~~00 = Non-combined. P0 is primary master, P2 is the primary slave, P1 is secondary master, P3 is the secondary slave.~~ ~~01 = Combined. IDE is primary. P1 is secondary master, P3 is the secondary slave.~~ ~~10 = Combined. P0 is primary master. P2 is primary slave. IDE is secondary~~ ~~11 = Reserved~~
~~7:6~~	~~SATA Mode Select: Software programs these bits to control the mode in which the SATA HBA should operate:~~ ~~00 IDE mode~~ ~~01 AHCI Mode~~ ~~10 RAID mode~~ ~~11 Reserved~~ ~~NOTES:~~ ~~1. When combined mode is used (non-zero MV), only IDE mode is allowed. IDE mode can be selected when RAID is enabled.~~ ~~2. AHCI mode may only be selected when MV=0. AHCI mode may be selected when RAID is enabled.~~ ~~3. RAID mode may only be selected when MV=0. Note that RAID5 mode may not be supported in certain SKUs of Intel® 631xESB/632xESB I/O Controller Hub.~~ ~~Programming these bits with values that are illegal (for example, selecting RAID when in combined mode) will result in-deterministic behavior by the hardware.~~

Bits	Description
5:2	Reserved.
1:0	Map Value — R/W. Map Value (MV): The value in the bits below indicate the address range the SATA ports responds to, and whether or not the PATA and SATA functions are combined. When in combined mode, the AHCI memory space is not available and AHCI may not be used. 00 = Non-combined. P0 is primary master, P2 is the primary slave, P1 is secondary master, P3 is the secondary slave. 01 = Combined. IDE is primary. P1 is secondary master, P3 is the secondary slave. 10 = Combined. P0 is primary master. P2 is primary slave. IDE is secondary 11 = Reserved
7:6	SATA Mode Select: Software programs these bits to control the mode in which the SATA HBA should operate: 00 IDE mode 01 AHCI Mode 10 RAID mode 11 Reserved NOTES: 1. When combined mode is used (non-zero MV), only IDE mode is allowed. IDE mode can be selected when RAID is enabled. 2. AHCI mode may only be selected when MV=0. AHCI mode may be selected when RAID is enabled. 3. RAID mode may only be selected when MV=0. Note that RAID5 mode may not be supported in certain SKUs of Intel® 631xESB/632xESB I/O Controller Hub. Programming these bits with values that are illegal (for example, selecting RAID when in combined mode) will result in deterministic behavior by the hardware.
5:2	Reserved.
1:0	Map Value — R/W. Map Value (MV): The value in the bits below indicate the address range the SATA ports responds to, and whether or not the PATA and SATA functions are combined. When in combined mode, the AHCI memory space is not available and AHCI may not be used. 00 = Non-combined. P0 is primary master, P2 is the primary slave, P1 is secondary master, P3 is the secondary slave. 01 = Combined. IDE is primary. P1 is secondary master, P3 is the secondary slave. 10 = Combined. P0 is primary master. P2 is primary slave. IDE is secondary 11 = Reserved

Bits	Description
7:6	SATA Mode Select: Software programs these bits to control the mode in which the SATA HBA should operate: 00 IDE mode 01 AHCI Mode 10 RAID mode 11 Reserved NOTES: 1. When combined mode is used (non-zero MV), only IDE mode is allowed. IDE mode can be selected when RAID is enabled. 2. AHCI mode may only be selected when MV=0. AHCI mode may be selected when RAID is enabled. 3. RAID mode may only be selected when MV=0. Note that RAID5 mode may not be supported in certain SKUs of Intel® 631xESB/632xESB I/O Controller Hub. Programming these bits with values that are illegal (for example, selecting RAID when in combined mode) will result in deterministic behavior by the hardware.
5:2	Reserved.
1:0	Map Value — R/W. Map Value (MV): The value in the bits below indicate the address range the SATA ports responds to, and whether or not the PATA and SATA functions are combined. When in combined mode, the AHCI memory space is not available and AHCI may not be used. 00 = Non-combined. P0 is primary master, P2 is the primary slave, P1 is secondary master, P3 is the secondary slave. 01 = Combined. IDE is primary. P1 is secondary master, P3 is the secondary slave. 10 = Combined. P0 is primary master. P2 is primary slave. IDE is secondary 11 = Reserved

Address Offset: 91h–93h
Default Value: 0000h
Attribute: R/W, R/WC, RO
Size: 24 bits

Bits	Description
23:21	Reserved.
20	**OOB Retry Mode**: 0 = The SATA controller will not retry after an OOB failure. 1 = The SATA controller will continue to retry after an OOB failure until successful (infinite retry).
19:16	Reserved.
15	**Port 3 Present (P3P)** — RO. Same as P0P, except for port 3.
14	**Port 2 Present (P2P)** — RO. Same as P0P, except for port 2.
13	**Port 1 Present (P1P)** — RO. Same as P0P, except for port 1.
12	**Port 0 Present (P0P)** — RO. When set, the SATA controller has detected the presence of a device on port 0. It may change at any time. This bit is cleared when the port is disabled via P0E. This bit is not cleared upon surprise removal of a device.

Bits	Description
11	**Port 3 Enabled (P3E)** — R/W. 0 = Disabled. The port is in the 'off' state and cannot detect any devices. 1 = Enabled. The port can transition between the on, partial, and slumber states and can detect devices. ***Note:*** This bit takes precedence over P2CMD.SUD (offset ABAR+298h:bit 1)
10	**Port 2 Enabled (P2E)** — R/W. 0 = Disabled. The port is in the 'off' state and cannot detect any devices. 1 = Enabled. The port can transition between the on, partial, and slumber states and can detect devices. ***Note:*** This bit takes precedence over P2CMD.SUD (offset ABAR+218h:bit 1)
9	**Port 1 Enabled (P1E)** — R/W. 0 = Disabled. The port is in the 'off' state and cannot detect any devices. 1 = Enabled. The port can transition between the on, partial, and slumber states and can detect devices. ***Note:*** This bit takes precedence over P1CMD.SUD (offset ABAR+198h:bit 1)
8	**Port 0 Enabled (P0E)** — R/W. 0 = Disabled. The port is in the 'off' state and cannot detect any devices. 1 = Enabled. The port can transition between the on, partial, and slumber states and can detect devices. ***Note:*** This bit takes precedence over P0CMD.SUD (offset ABAR+118h:bit 1)
7:4	Reserved
3	**Port 5Present (**P5P**)** — RO. When set, the SATA controller has detected the presence of a device on port 5. It may change at any time. This bit is cleared when the port is disabled via P5E. This bit is not cleared upon surprise removal of a device.
2	**Port 4Present (P4P)** — RO. When set, the SATA controller has detected the presence of a device on port 4. It may change at any time. This bit is cleared when the port is disabled via P4E. This bit is not cleared upon surprise removal of a device.
1	**Port 5Enabled (P5E)** — R/W. The definition of this bit is the same as P0E, except for port 5. This bit takes precedence over P5CMD.SUD.
0	**Port 4 Enabled (P4E)** — R/W. The definition of this bit is the same as P0E, except for port 4. This bit takes precedence over P4CMD.SUD.

10. Content updated to Table6-1 and Table6-2

Table6-1 Intel® 631xESB/632xESB I/O Controller Hub Power ~~Consumption Estimates~~ Requirements with Wake-on-LAN(LAN tied to AUX power)

Platform Voltage Rail	Current in S0	Current in S3	Current in S4	Current in S5	Current in G3
Logic 1.5V Core	3.82 A	N/A	N/A	N/A	N/A
I/O 1.5V Core	1.98A	N/A	N/A	N/A	N/A
1.5V Aux	1.97A	0.62A	~~N/A~~ 0.62A	~~N/A~~ 0.62A	N/A
3.3V Core	1.8A	N/A	N/A	N/A	N/A
3.3V Sus	70mA	30mA	40mA	40mA	N/A
3.3V Aux	0.27A	0.07A	N/A	N/A	N/A
Vcc5REF	150uA	N/A	N/A	N/A	N/A
Vcc5REFSus	10mA	10mA	10mA	10mA	OFF
VccRTC	N/A	N/A	N/A	N/A	6uA

Table6-2 Intel® 631xESB/632xESB I/O Controller Hub Power ~~Consumption Estimates with~~ Requirements without Wake-on-LAN(LAN tied to CORE power)

Platform Voltage Rail	Current in S0	Current in S3	Current in S4	Current in S5	Current in G3
Logic 1.5V Core	5.54A	N/A	N/A	N/A	N/A
I/O 1.5V Core	2.23A	N/A	N/A	N/A	N/A
1.5V Aux	~~2.07~~ N/A	N/A	N/A	N/A	N/A
3.3V Core	~~N/A~~ 2.07A	N/A	N/A	N/A	N/A
3.3V Sus	70mA	30mA	40mA	40mA	N/A
3.3V Aux	N/A	N/A	N/A	N/A	N/A
Vcc5REF	150uA	N/A	N/A	N/A	N/A
Vcc5REFSus	10mA	10mA	10mA	10mA	~~N/A~~ OFF
VccRTC	N/A	N/A	N/A	N/A	6uA

Notes:
1: These numbers are pre-silicon estimates. Actual numbers will be included after measured.

11. Content updated to Table6-8

Table6-8 Intel® 631xESB/632xESB I/O Controller Hub 1.5 V Supply Rail Tolerances

Symbol	Parameter	AC Min	DC Min	Nom	DC Max	ACMax	Unit	Notes
Vcc, VccSUS[2], ~~VREFPCI~~	1.5 V Supply Voltage	1.425	1.455	1.5	1.575	1.605	V	Note 1&2
VREFPCI	Analog Reference Voltage to PCI	0.728		0.75	0.773		V	

12. Content updated to Table2-32

Table2-32 Functional Strap Definitions

STRAP_4	Reserved	N/A	**Strap Pin.** ~~pull up to~~ Connected to P1V5_ESB.
STRAP_5	Reserved	N/A	**Strap Pin.** ~~pull up to~~ Connected to P1V5_ESB.
STRAP_6	Reserved	N/A	**Strap Pin.** ~~pull up to~~ Connected to P1V5_AUX.

13. Content updated to Section13.6.1.64

Section 13.6.1.64 Offset DCh: PX_BSTS—PCI-X* Bridge Status Register (Bm:D0:F3)

Bits	Type	Reset	Description
15:8	~~RW~~ RO	0	~~Reserved.~~ Bus number.
7:3	RO	0	~~Reserved.~~ Device number.

14. Clear figures updated to Figure 9-1 to Figure 9-4

Figure 9-1 Mechanical Layout

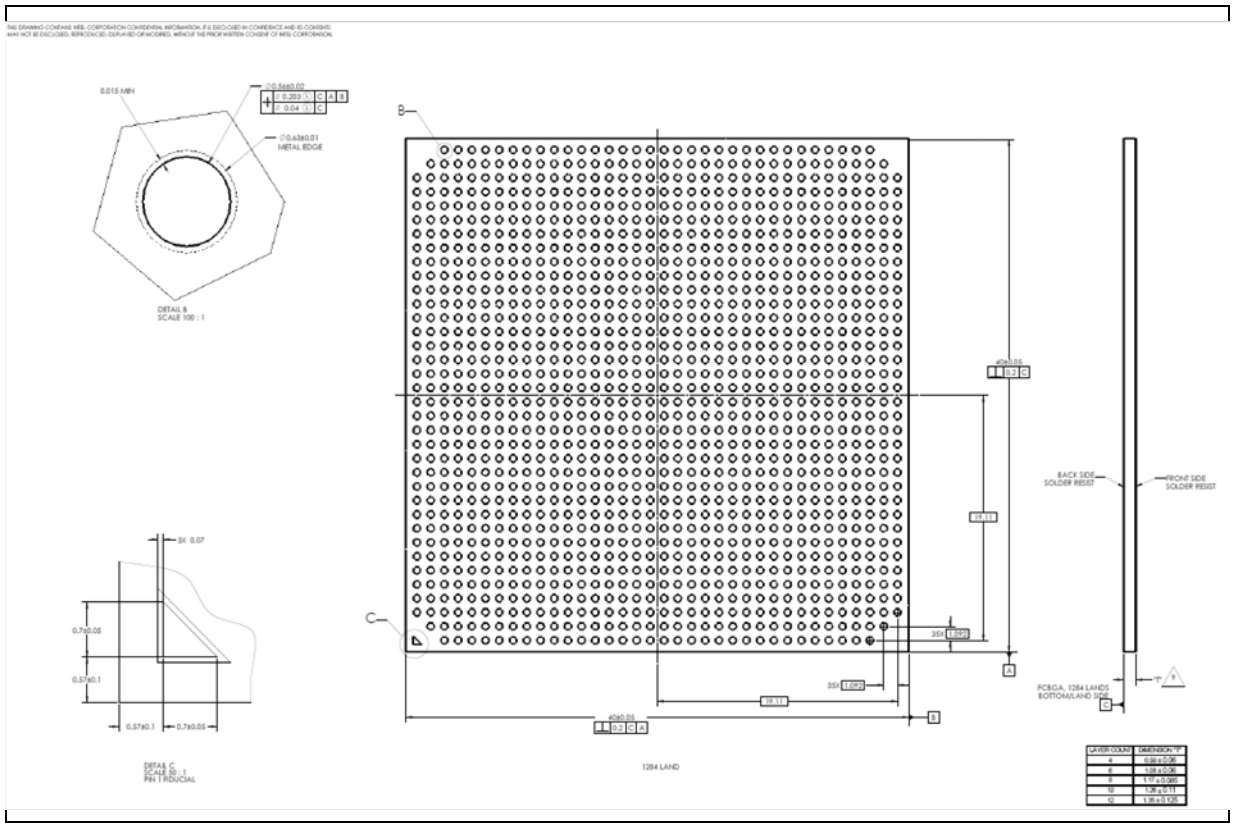

Figure 9-2 Mechanical Layout

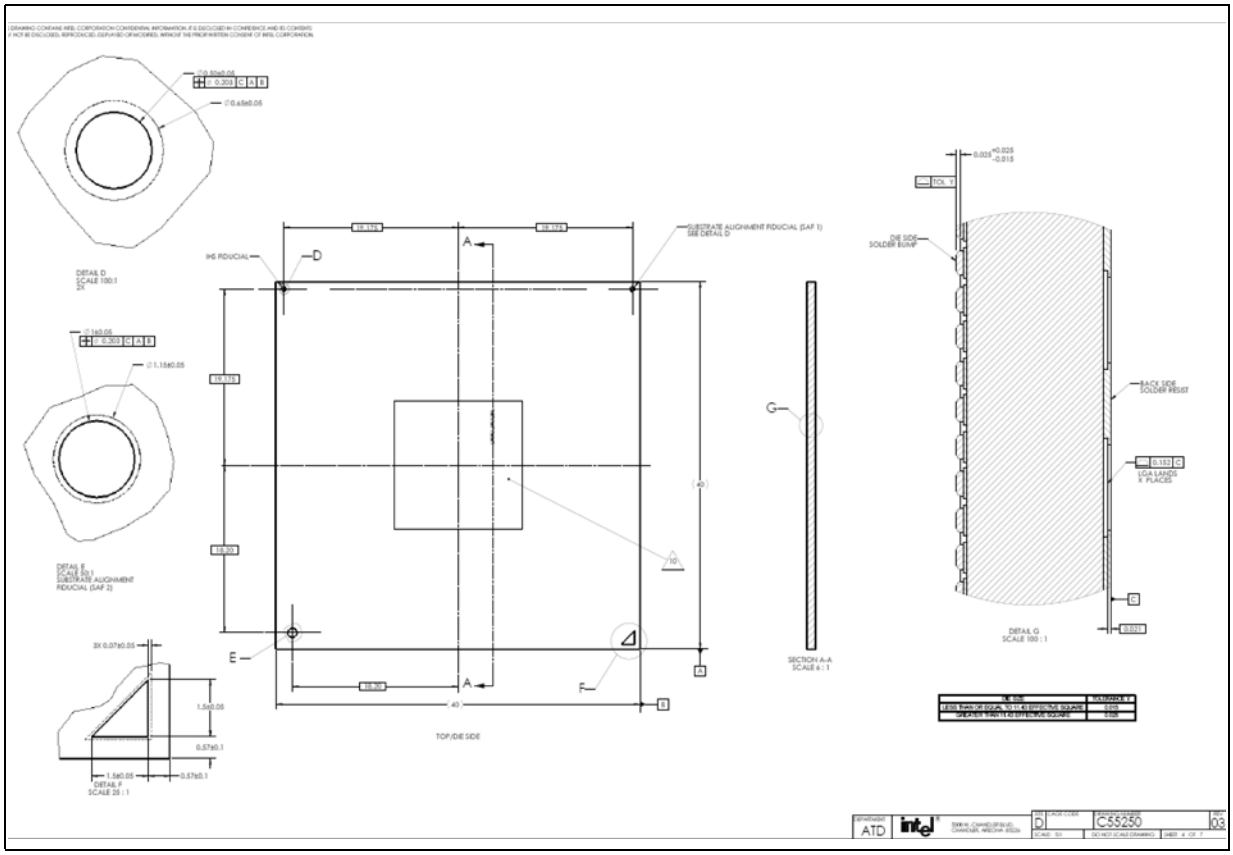

Figure 9-3 Mechanical Layout

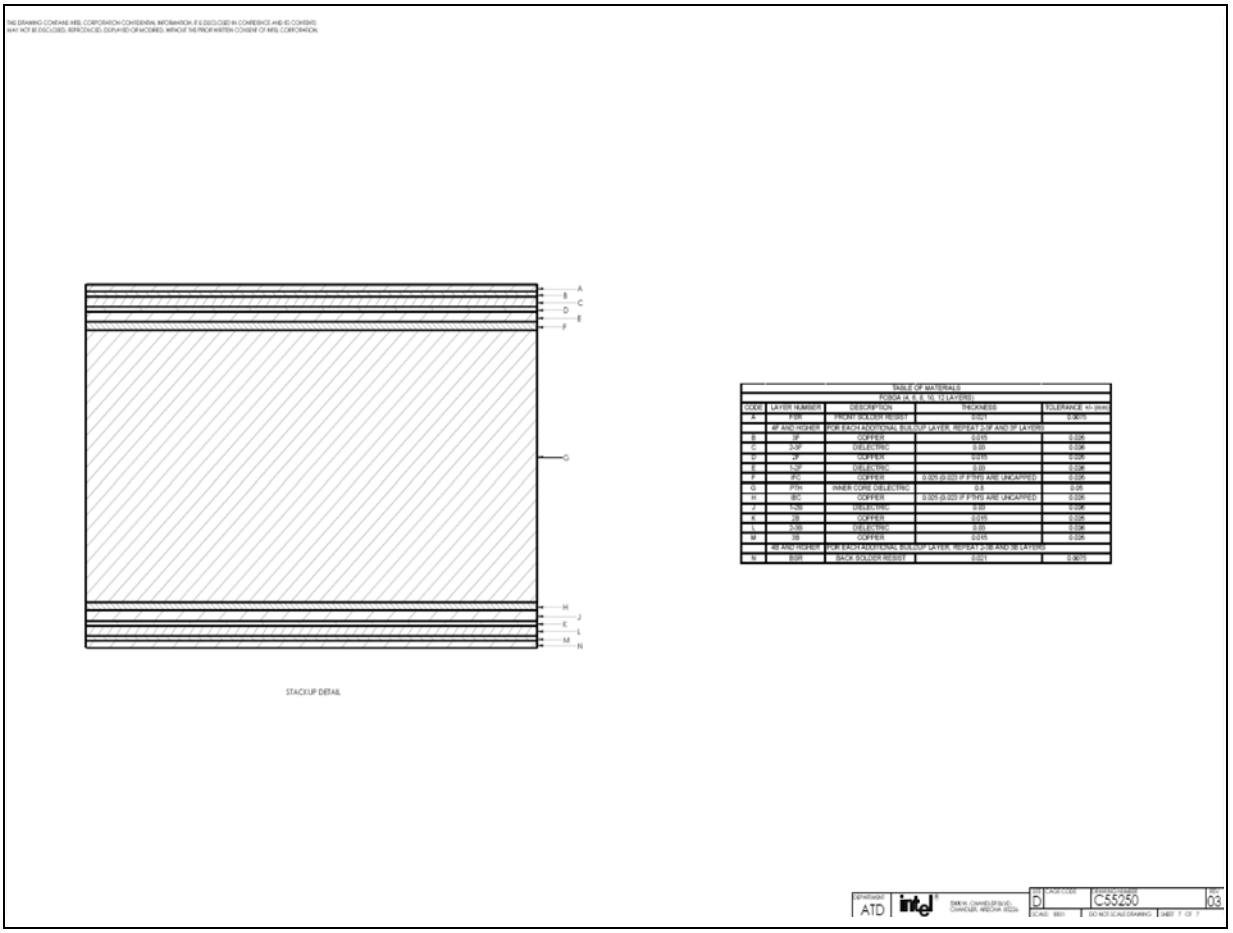

Figure 9-4 Mechanical Layout

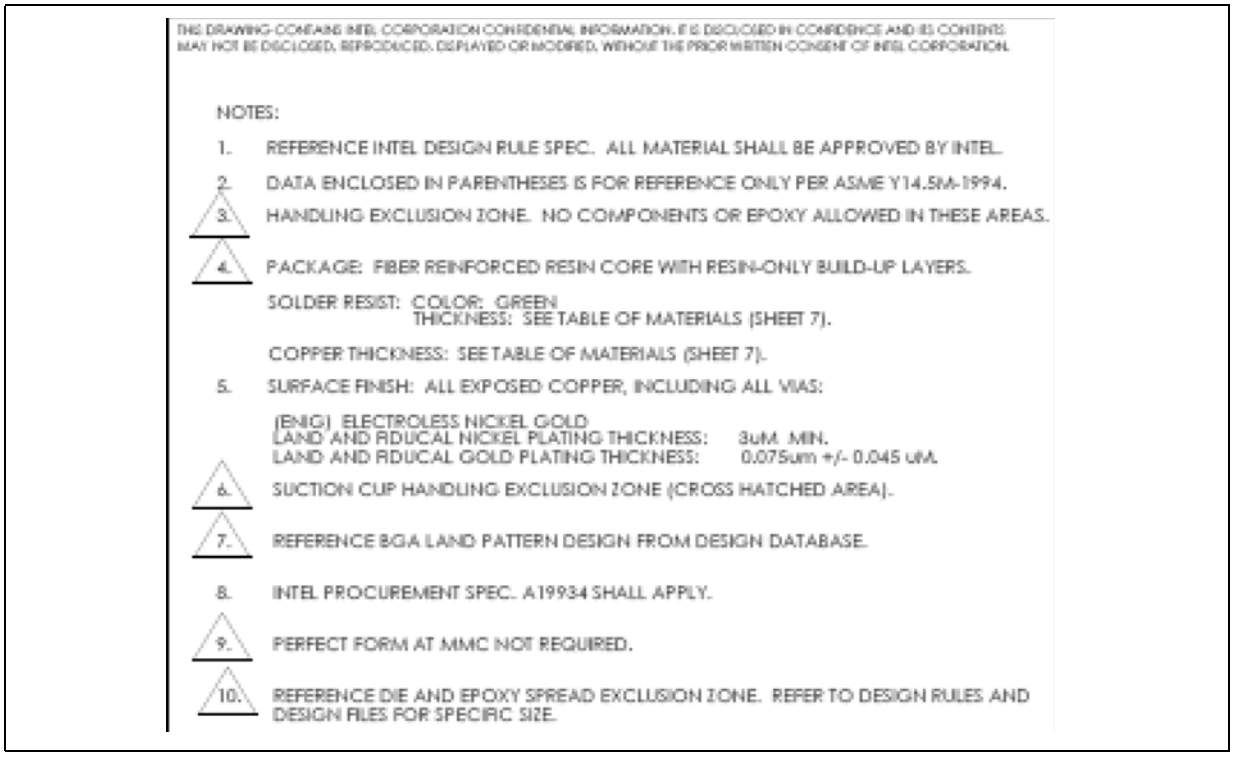

15. Added Table6-20 Intel® 631xESB/632xESB I/O Controller Hub Absolute Maximum Ratings

Table 6-20 Intel® 631xESB/632xESB I/O Controller Hub Absolute Maximum Ratings

Parameter	Maximum Limits (V)
5.0V Supply Voltage with respect to VSS	-0.5 to 5.5
3.3V Supply Voltage with respect to VSS	-0.5 to 4.0
1.5V Supply Voltage with respect to VSS	-0.5 to 2.1
Vtt Supply Voltage with respect to VSS	-0.5 to 2.1

16. Section 8.1, changed pin AA36 name from “Reserved” to “PHY_POWER_DOWN”, changed pin AB36 name from “Reserved” to “PHYRST#”

17. Section 8.2, changed pin AA36 name from “Reserved” to “PHY_POWER_DOWN”, changed pin AB36 name from “Reserved” to “PHYRST#”

18. Input Signal Behavior Clarifications

In Section 4.3 and 4.4, a clarification is added to indicate that input signals may act as outputsduring indeterminate state as indicated below:

"Intel® 631xESB/632xESB I/O Controller Hub suspend well signal states are indeterminate and undefined and may glitch, including input signals acting as outputs prior to RSMRST# de-assertion. This does not apply to LAN_PWR_GOOD, SLP_S3#, SLP_S4# and SLP_S5#. These signals are determinate and defined prior to RSMRST# de-assertion.

Intel® 631xESB/632xESB I/O Controller Hub core well signal states are indeterminate and undefined and may glitch, including input signals acting as outputs prior to PWROK assertion. This does not apply to FERR# and THRMTRIP#. These signals are determinate and defined prior to PWROK assertion."

19. APM_CNT and APM_STS Registers

In Section 21.8 for Power Management, add Section "APM I/O Decode"

APM I/O Decode

Table below shows the I/O registers associated with APM support. This register space is enabled in the PCI Device 31: Function 0 space (APMDEC_EN), and cannot be moved (fixed I/O location).

APM Register Map

Address	Mnemonic	Register Name	Default	Type
B2h	APM_CNT	Advanced Power Management Control Port	00h	R/W
B3h	APM_STS	Advanced Power Management Status Port	00h	R/W

APM_CNT—Advanced Power Management Control Port Register

I/O Address: B2h
Default Value: 00h
Lockable: No
Power Well: Core
Attribute: R/W
Size: 8-bit
Usage: Legacy Only

Bit	Description
7:0	Used to pass an APM command between the OS and the SMI handler. Writes to this port not only store data in the APMC register, but also generates an SMI# when the APMC_EN bit is set.

APM_STS—Advanced Power Management Status Port Register

I/O Address: B3h
Default Value: 00h
Lockable: No
Power Well: Core
Attribute: R/W
Size: 8-bit
Usage: Legacy Only

Bit	Description
7:0	Used to pass data between the OS and the SMI handler. Basically, this is a scratchpad register and is not affected by any other register or function (other than a PCI reset).